Imperial College of Science, Technology and Medicine
Department of Physics, Centre for Complexity Science

Studies in Complex Systems: Complex Networks with Hidden Layers & Modelling Atrial Fibrillation

Max Falkenberg

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Physics of Imperial College, October 2021

Statement of originality

I declare that the work presented in this thesis is entirely my own work unless stated otherwise. I have acknowledged contributions from collaborators and others according to the standard referencing practices. I have sought permission to reproduce any third-party copyrighted material in this thesis from the copyright holders and collaborators. Permissions are included at the end of this thesis.

Max Falkenberg, October 2021

Copyright declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

This thesis, covering the broad theme of complexity science, is split into two parts.

In part 1, I discuss models of network growth with hidden network layers. First, I introduce the "k2 model", based on linear preferential attachment, where the attachment probability is proportional to the number of nodes within two steps of a target node. The model generates time dependent degree distributions and attachment kernels, despite initially appearing to grow as linear preferential attachment. Subsequently, I focus on node copying models. Motivated by observations of asymmetric triadic closure in real networks, I introduce a model of heterogeneous copying, built using a hidden and observed network layer. Framed in a social context, these two layers represent a node's inner social circle, and wider social circle, such that the model can bias copying probabilities towards, or against, a node's inner circle of friends. Compared to homogeneous copying, this heterogeneous model suppresses power-laws, leads to unusually high clustering, and results in size-independent clique growth.

Part 2 discusses simple percolation-based models of atrial fibrillation (AF), the most common cardiac arrhythmia. First, I introduce a 3d-analogue of a previous 2d lattice model. The model highlights the unification of previous discrepant observations in the clinical literature, specifically, how asymmetry in the fibre structure between the inside and the outside of the heart leads to differences in the observed surface activation patterns. The 3d model is not representative of real atrial geometries. Hence, I subsequently introduce a further adaptation of the model in which signals propagate on a spatial network representation of the atria, which can be adapted for use with real imaging data. The model highlights how, as the atria become increasingly decoupled through the accumulation of interstitial fibrosis, the successful termination of fibrillation via ablation becomes increasingly difficult.

Publications

The following publications form the basis of this thesis:

[1] M. Falkenberg *et al.* "Identifying time dependence in network growth". In: *Physical Review Research* 2.2 (2020), p. 023352. DOI: 10.1103/PhysRevResearch.2.023352

[2] M. Falkenberg. "Heterogeneous node copying from hidden network structure". In: *Communications Physics* 4.1 (2021), p. 200. DOI: 10.1038/s42005-021-00694-1

[3] M. Falkenberg *et al.* "Unified mechanism of local drivers in a percolation model of atrial fibrillation". In: *Physical Review E* 100.6 (2019), p. 062406. DOI: 10.1103/PhysRevE.100.062406

[4] M. Falkenberg *et al.* "Identifying Potential Re-Entrant Circuit Locations From Atrial Fibre Maps". In: *2019 Computing in Cardiology (CinC)*. IEEE. (2019), pp. 1–4. DOI: 10.22489/CinC.2019.102

[5] M. Falkenberg *et al.* "Identifying locations susceptible to micro-anatomical reentry using a spatial network representation of atrial fibre maps". In: *bioRxiv* (2021). DOI: 10.1101/2021.09.13.460069

Publications completed during my PhD which are not included in this thesis:

[6] M. Falkenberg *et al.* "Machine learning methods for locating re-entrant drivers from electrograms in a model of atrial fibrillation". In: *Royal Society Open Science* 5.4 (2018), p. 172434. DOI: 10.1098/rsos.172434

[7] N. R. Franks *et al.* "Digging the optimum pit: antlions, spirals and spontaneous stratification". In: *Proceedings of the Royal Society B* 286.1899 (2019), p. 20190365. DOI: 10.1098/rspb.2019.0365

[8] A. Ciacci *et al.* "Understanding the transition from paroxysmal to persistent atrial fibrillation". In: *Physical Review Research* 2.2 (2020), p. 023311. DOI: 10.1103/PhysRev

Research.2.023311

[9] H. Rajpal *et al.* "Interpretable XGBoost Based Classification of 12-lead ECGs Applying Information Theory Measures From Neuroscience". In: *2020 Computing in Cardiology*. IEEE. (2020), pp. 1–4. DOI: 10.22489/CinC.2020.185

Acknowledgements

At the start of my third year of physics undergrad, I had next to no intention of doing a PhD, particularly given my dislike for attention to detail and my, at best, mediocre maths skills. That year, I was lucky enough to take the Statistical Mechanics course led by Matthew Foulkes and Kim Christensen. A term later I took the, largely computational, Complexity & Networks course, led by Tim Evans and, again, Kim. It is because of those two courses that I am where I am today, and for that I will be eternally grateful.

It is perhaps no surprise that, given his role in the two courses which inspired me to do a PhD, I decided to a summer project with Kim after my third year, followed by my Master's project in fourth year. Since then, I have been fortunate to have Kim as my ever gracious PhD supervisor guiding me through the chaos which is academia. Thank you, Kim, for your kindness, support, motivation and drive, which has seen me through this PhD – you have been a star.

Although I don't think I ever officially had a second supervisor, in practice that role was taken by Tim Evans, to whom I am also extremely grateful, especially for his willingness to put up with my constant badgering to do paperwork for me. Without my many conversations with Tim, I am almost certain that a fair chunk of my work on networks would never have happened. Thank you.

I must also thank the many other academics who have played a role in my PhD: Nigel Franks and Ana-Sendova Franks who showed me the joys of insects and the fun in working on a project completely removed from your main research. Also to Nigel for laughing at my self-imposed deadlines, and for teaching me that the formula to accurate time estimates is triple your initial guess plus three months. Nicholas Peters for showing me the glorious world of cardiology. Fu Siong Ng for engaging with me and my work even though he was under no obligation to do so, and to Henrik Jeldtoft Jensen for his support in his role running the Centre for Complexity Science.

Of course, no one survives a PhD without their fellow PhDs and postdocs who make the day-to-day experience worth it. For me, those people were Hardik, Vaiva, Fernando, Pedro, Mada, Alberto, Chester, and others too numerous to list. Thank you for making the last few years a lot of fun. Thank you also to my co-authors – there have been many, and without them, most of the work in this thesis would not exist. The list is long so I will not thank everyone

individually, but special credit should go to James Coleman, Jong-Hyeok Lee and Ken-Ichiro Ogawa who all played critical roles in the projects they were involved in and to whom I am indebted.

Finally, to my friends outside of work, Sam, Josh, Tim, Jake, Ollie, Eden, Mada, Andrei and all the others I should thank but have forgotten to list. To my Mum and Dad for being supportive and always letting me do what I wanted to do. To Alex for being Alex. To Flo, Lennox, Lola, Nora and Oscar for adding a dose of chaos to my life. And of course to my darling Becky, for putting up with me, supporting me, proof reading an awful lot of bad work, drawing me astoundingly good diagrams on request, and aggressively cutting the word 'hence' from every second sentence where I've used it. I couldn't have done it without you.

If you do choose to read this thesis in full – which, in all honesty, I'm not sure I would recommend – I hope you enjoy it.

Max Falkenberg, October 2021

A note to the reader

This thesis is made up of two parts, the first focusing on complex networks with hidden layers, and the second focusing on the modelling of atrial fibrillation. Within each part, there are two principle research chapters, each of which is largely stand alone. As a result, I have decided that, rather than include a single introduction for the whole thesis, each part starts with a brief overview and motivation, with a more comprehensive introduction at the start of each self-contained chapter. This thesis also does not include any appendices – as an alternative, I have included any additional technical details or results at the end of each chapter. Any material presented after the chapter summary can be treated as supplementary and can be skipped if the reader so wishes. Finally, the reader may notice that, particularly in my overviews and conclusions, I have taken the liberty to express more of my own opinion than would usually be present in a published paper – I do hope the reader will forgive this lack of formality, after all, it is not every day I get to write my one and only thesis.

Dedication

Für meine Oma. Das hätte sie gefreut!

Contents

Abstract i

Publications ii

Acknowledgements iv

Symbols & acronyms for part 1, chapters 1–3 xx

Symbols & acronyms for part 2, chapters 4–6 xxiv

Part 1 : Complex networks with hidden layers 2

1 Overview & motivation 2

2 Identifying time dependence in network growth 5

2.1 Publication 5

2.2 Introduction 6

2.3 Methods 12

2.3.1 Model definition 12

2.3.2 Measuring the time dependence of preferential attachment 16

2.4 Results 17

2.4.1 Simulation results 17

2.4.2 Mathematical results 23

2.4.3 Application to real world networks 29

2.5 Discussion 33

2.6 Conclusions 36

2.7 Chapter summary 37

2.8 Mathematical derivations 37

2.8.1 Derivation of Eq. (2.1). 37

2.8.2 Derivation of Eq. (2.10). 39

3 Heterogeneous node copying from hidden network structure 41

3.1 Publication 41

3.2 Introduction 41

3.3 Correlated copying model 45

3.3.1 Basic topological properties 46

3.3.2 Degree distribution 47

3.3.3 Clique distribution 51

3.3.4 Clustering 56

3.3.5 Path lengths 58

3.4 General correlated copying model 59

3.5 Comparing copying models 64

3.6 Discussion 65

3.7 Conclusions 68

3.8 Chapter summary 69

3.9 Mathematical derivations 70

3.9.1 Observed network second moment 70

3.9.2 Observed degree distribution 72

3.9.3 Derivation of Eq. (3.9) 72

Part 2 : Modelling atrial fibrillation **76**

4 Overview & motivation **76**

5 A percolation model of atrial fibrillation **82**

5.1 Publication 82

5.2 Introduction 82

5.3 Model definition & properties 84

5.4 Results 88

5.5 Discussion 92

5.6 Limitations & next steps 94

5.7 Chapter summary 95

5.8 Technical details 96

5.8.1 Model definition 96

5.9 Simulation details 99

5.9.1 Ablation methods 99

5.10 Further model results 102

5.10.1 Depth of reentrant circuits 102

6 Identifying micro-anatomical reentry using spatial networks 105

6.1 Publications 105

6.2 Introduction 106

6.3 Aims 107

6.4 Materials and methods 109

6.4.1 Atrial datasets 110

6.4.2 Spatial network with fibre structure 113

6.4.3 Spatial network without fibre structure 119

6.4.4 Analysis methods 121

6.5 Results 123

6.5.1 Spatial distribution of identified micro-reentrant substrate 123

6.5.2 Spatial clustering of micro-reentrant substrate 127

6.5.3 Occupied voxel fraction 128

6.5.4 Comparing fibre and fibre-less micro-reentrant substrates 129

6.5.5 Comparing atrial datasets 133

6.6 Discussion 135

6.6.1 Limitations . . . 138

6.7 Conclusions . . . 139

6.8 Chapter summary . . . 141

6.9 Technical details . . . 142

6.9.1 Fibre tractography . . . 142

6.9.2 Discrete diffusion model . . . 145

6.9.3 Atrial wall thickness & occupied voxel fraction . . . 148

6.10 Supplementary results . . . 152

6.10.1 Spatial distribution of risk substrate in the sheep atria . . . 152

6.10.2 Effect of fibre structure in sheep atria . . . 160

6.10.3 Statistical analysis of OVF and LCF distributions . . . 162

Bibliography **165**

Right to reuse figures **197**

List of Tables

3.1 A comparison between the correlated copying model and the uniform copying model. 65

3.2 Clustering values for a selection of sparse undirected networks which may plausibly grow via a copying mechanism. 66

6.1 Summary information for the three image-based atrial fibre orientation datasets. 111

6.2 Spatial network coupling and risk values for the fibre model. 115

6.3 Spatial network coupling and risk values for the fibre-less null model. 120

6.4 Summary of the differences between fibre tractography using FACT and ESS. . . 143

6.6 Summary of the occupied voxel fraction statistical analysis for the healthy sheep atria. 165

6.7 Summary of the occupied voxel fraction statistical analysis for the heart failure sheep atria. 165

6.5 Summary of the occupied voxel fraction statistical analysis for the human atria. 165

6.8 Summary of the longitudinal connection fraction statistical analysis for the human atria. 166

6.9 Summary of the longitudinal connection fraction statistical analysis for the healthy sheep atria. 166

6.10 Summary of the longitudinal connection fraction statistical analysis for the heart failure sheep atria. 166

List of Figures

2.1 A graph schematic indicating the difference between the first and second degree of a node. 13

2.2 A schematic illustrating how the k2 model can be thought of as an observed and influence network. 15

2.3 The degree distribution and relative attachment kernels for the k2 model. 18

2.4 An illustration of the time dependence of the relative attachment kernel in the k2 model. 21

2.5 An illustration of the time independence of the relative attachment kernel for the BA model. 22

2.6 The growth in the number of edges in the k2 model's influence network. 27

2.7 The degree evolution of individual nodes in the k2 model. 28

2.8 The ratio of the cumulative relative attachment kernels for six real world networks. 30

2.9 A figure demonstrating the equivalence between the sum over all nodes of the second degree and of the first degree of each node squared. 39

3.1 Two models of network formation via copying: the uniform copying model and the correlated copying model. 45

3.2 The degree distributions for the correlated copying model and uniform copying model. 50

3.3 A schematic illustrating the number of triangles formed in a single time step of the correlated copying model. 52

3.4 The average rescaled number of cliques of size n in the observed network. 55

3.5 Individual trajectories for the rescaled number of cliques of size n as a function of time. 56

3.6 Properties of the general correlated copying model. 60

3.7 The observed degree distributions for the general correlated copying model. . . . 63

4.1 A schematic showing the four heart chambers. 76

4.2 A schematic showing the heart's electrical system. 78

4.3 An illustration showing the ablation of drivers in atrial fibrillation. 79

4.4 A schematic showing the formation of a micro-anatomical reentrant circuit. . . . 80

5.1 The emergence of atrial fibrillation in the 3d percolation model. 86

5.2 A lattice cross-section showing the presence of a reentrant circuit. 88

5.3 The average time spent in AF for the homogeneous and inhomogeneous models. 89

5.4 Activation maps showing the presence of a reentrant circuit observed experimentally and in the percolation model. 90

5.5 The probability of successfully ablating a reentrant circuit. 91

5.6 Activation maps demonstrating the failure of partial depth ablation. 93

5.7 A schematic outlining the rules of the cellular automata. 98

5.8 The formation of a reentrant circuit in the cellular automata model. 100

5.9 The distribution of reentrant circuit depths in the homogeneous model. 103

5.10 The distribution of reentrant circuit depths in the inhomogeneous model. 104

6.1 A schematic illustrating how we model the accumulation of fibrosis using spatial networks. 108

6.2 A schematic illustrating the construction of spatial networks from imaging data. 109

6.3 A schematic showing how local fibre orientation vectors are converted into a spatial network. 113

6.4 A schematic showing how the null model spatial networks are constructed. . . . 119

6.5 A schematic outlining the definition of the occupied voxel fraction. 123

6.6 The identified micro-reentrant substrate for the human fibre model. 124

6.7 The identified micro-reentrant substrate for the human fibre-less null model. . . 125

6.8 A figure illustrating how the micro-reentrant substrate becomes less clustered as the risk of micro-reentry increases. 128

6.9 A figure illustrating that the micro-reentrant substrate is biased to atrial regions which are thin when the risk of micro-reentry is low. 129

6.10 An illustration of the competing factors determining which regions are susceptible to micro-anatomical reentry. 130

6.11 A figure illustrating how the micro-reentrant substrate correlates with atrial thickness and fibre alignment in the human atria. 132

6.12 Heatmaps illustrating node density fluctuations from two tractography methods. 143

6.13 A schematic illustrating the mechanism for detecting a micro-reentrant circuit. . 147

6.14 Plots of the atrial thickness for each dataset. 149

6.15 The occupied voxel fraction for each dataset. 150

6.16 Atrial cross-sections for each dataset showing the occupied voxel fraction. 151

6.17 A comparison between the occupied voxel fraction and wall thickness. 151

6.18 The spatial distribution of the micro-reentrant substrate in the fibre model for the healthy sheep atria. 154

6.19 The spatial distribution of the micro-reentrant substrate in the null model for the healthy sheep atria. 155

6.20 Additional views for the spatial distribution of the micro-reentrant substrate in the healthy sheep atria. 156

6.21 The spatial distribution of the micro-reentrant substrate in the fibre model for the heart failure sheep atria. 157

6.22 The spatial distribution of the micro-reentrant substrate in the null model for the heart failure sheep atria. 158

6.23 Additional views for the spatial distribution of the micro-reentrant substrate in the heart failure sheep. 159

6.24 A figure illustrating how the micro-reentrant substrate correlates with atrial thickness and fibre alignment in the healthy sheep atria. 160

6.25 A figure illustrating how the micro-reentrant substrate correlates with atrial thickness and fibre alignment in the heart failure sheep atria. 161

Symbols & acronyms for part 1, chapters 1–3

$a_{\alpha,\beta}$	The adjacency matrix encoding edges between nodes α and β
α	Node label (usually the newly added node)
BA	Barabási-Albert
β	Node label (usually the target node for copying)
$cc(\alpha)$	The clustering coefficient for node α
$CC(G_O)$	The global clustering coefficient for the observed network
CCM	Correlated copying model
δ	Degree evolution exponent assuming preferential attachment (linear or non-linear)
$E^{(2)}(t)$	Number of edges in the influence network at time t
E_H	The set of edges in the hidden network
$E_H(t)$	The number of edges in the hidden network at time t
E_O	The set of edges in the observed network
$E_O(t)$	The number of edges in the observed network at time t
ε	Densification exponent
$\langle\cdot\rangle$	Expectation value
$\langle\cdot \mid \cdot\rangle$	Conditional expectation value
$f(\cdot)$	Numerator of the preferential attachment kernel
γ	Exponent for power-law preferential attachment kernel
γ_j	Node label for the neighbours of the target node β
G	The graph G

GCCM	Generalised correlated copying model
G_H	The hidden network
G_O	The observed network
H_n	The n^{th} harmonic number
$H_n^{(m)}$	The n^{th} generalised harmonic number of order m
k	Node degree (equivalent to $k^{(1)}$)
k_H	Node degree in the hidden network
$k^{(1)}$	First degree of a node (equivalent to k)
$k_i^{(1)}$	First degree of node i specifically
$k^{(2)}$	Second degree of a node
$k_i^{(2)}$	Second degree of node i specifically
$k_{\text{hub}}^{(1)}$	The observed degree of a hub node in the k2 model
k_O	Node degree in the observed network
$\tilde{k}_O$	Initial node degree in the observed network
κ	Power-law degree distribution exponent
ℓ	Shortest path length
$(\ell)_{ij}$	Shortest path length between nodes i and j
ℓ_H	Shortest path length in hidden network
$\langle L_H(t) \rangle$	Average shortest path length in hidden network
ℓ_O	Shortest path length in observed network
$\langle L_O(t) \rangle$	Average shortest path length in observed network

m	Number of edges added to the network at each time step
m_0	Number of nodes in the initial network
$n_\ell(t)$	Number of times node ℓ appears in the second degree summation at time t
$N(t)$	Number of nodes in the network at time t
p	Probability
p_{eff}	The effective copying probability
p_{hid}	The copying probability for edges in the hidden network (inner copying)
p_{obs}	The copying probability for edges in the observed network only (outer copying)
$P(k)$	Degree distribution
$\mathcal{P}(\ell;t)$	Distribution of shortest path lengths at time t
$\mathcal{P}_H(\ell_H;t)$	Distribution of shortest path lengths in the hidden network at time t
$\mathcal{P}_O(\ell_O;t)$	Distribution of shortest path lengths in the observed network at time t
$p_H(k_H)$	Degree distribution for the hidden network
$p_\infty(k)$	Degree distribution as a function of node degree in the thermodynamic limit
$p_O(k_O)$	Degree distribution for the observed network
$\tilde{p}_O(k_O)$	Cumulative degree distribution for the observed network
$\phi_t(k,k')$	Relative attachment kernel between nodes of degree k and k'
$\tilde{\phi}_t(k,k')$	Cumulative relative attachment kernel between nodes of degree k and k'
$\pi_O(k_O)$	Contribution to the degree distribution from nodes in the pre-existing network
$\Pi(k)$	Preferential attachment kernel as a function of node degree
$\Pi^{(BA)}$	Preferential attachment kernel of the BA model

$\Pi_i^{(BA)}$	Preferential attachment kernel of node i in the BA model
$\Pi^{(k2)}$	Preferential attachment kernel of the k2 model
$\Pi_i^{(k2)}$	Preferential attachment kernel of node i in the k2 model
$Q_n(t)$	The number of cliques of size n at time t
$Q_n^D(t)$	Cliques of size n from direct attachment at time t
$Q_n^I(t)$	Cliques of size n from induced attachment at time t
$S_{(1)}(t)$	Sum over first degree squared of all nodes in the network at time t
$S_{(2)}(t)$	Sum over second degree of all nodes in the network at time t
$s_2(t)$	The number of star graphs of size 2 at time t
$\sigma^2(k_O)$	Degree variance in observed network
t	Time in arbitrary units
t_α	Time node α was added to the network
τ_{G_O}	Observed network transitivity
UCM	Uniform copying model
V	The set of nodes in graph G
$\xi_i(t)$	Contribution to the second degree of node i from non-neighbouring nodes at time t

Symbols & acronyms for part 2, chapters 4–6

AF	Atrial fibrillation
b	Course graining factor
c	Characteristic coupling distance
CT	Crista terminalis
CS	Coronary sinus
δ	Fraction of nodes susceptible to conduction block
DDM	Discrete diffusion model
$\Delta\theta$	Average fibre angle
$\Delta\theta_{\mathrm{ENDO}}$	Average fibre angle on endocardium
$\Delta\theta_{\mathrm{EPI}}$	Average fibre angle on epicardium
d_{sep}	Minimum node separation distance
ϵ	Probability of conduction block
ESS	Evenly spaced streamlines
FO	Fossa ovalis
FOC	Fibre orientation correlation
H	High risk
HF	Heart failure
IAS	Inter-atrial septum
IVC	Inferior vena cava
ℓ	Re-entrant pathway length

L	Low risk
λ	Risk parameter (equivalent to the rate parameter in a Poisson process)
LGE-MRI	Late-gadolinium enhance magnetic resonance imaging
LAA	Left atrial appendage
LCF	Longitudinal connection fraction
LIPV	Left inferior pulmonary vein
LSPV	Left superior pulmonary vein
L_x	Lattice width
L_y	Lattice length
L_z	Lattice depth
M	Medium risk
MFA	Multiple focal ablations
MRI	Magnetic resonance imaging
MV	Mitral valve
N	Number of re-entrant circuits
N_v	Number of re-entrant circuits detected in voxel v
$\bar{\nu}$	Average lattice coupling probability
$\nu_{\parallel}$	Longitudinal coupling probability on lattice
$\nu_{\perp}$	Transverse coupling probability on lattice
OVF	Occupied voxel fraction
PLA	Posterior left atrium

PM	Pectinate muscles
PRA	Posterior right atrium
$p(x; r, c)$	Transverse connection probability in spatial network
PV	Pulmonary vein
r	Arbitrary steepness parameter
$\tilde{r}$	OVF calculation radius
RAA	Right atrial appendage
RIPV	Right inferior pulmonary vein
RSPV	Right superior pulmonary vein
$\tilde{R}_v$	Rescaled rate of reentry in voxel v
SFA	Single focal ablation
SVC	Superior vena cava
t	Time
T	Pacing frequency
τ	Refractory period
$\tilde{\tau}$	Refractory wavelength
t_i	Time at which nodes enter the refractory period
TV	Tricuspid valve
v	Voxel label
x	Distance separating two nodes in spatial network

Part 1

Complex networks with hidden layers

Chapter 1

Overview & motivation

The study of complex networks has expanded rapidly over the past 20 years. Many real systems have been analysed using networks with great success, showing many non-trivial properties [10]. Model networks have been defined to understand the origin and development of these properties from elementary principles. For instance, the Watts-Strogatz model generates networks with short average path lengths but high clustering coefficients, explaining the small world phenomenon [11]. Similarly, the Barabási-Albert (BA) model, an undirected version of the Price model [12], demonstrates that scale free (power-law) degree distributions in real networks can arise from a combination of growth and preferential attachment [13]. These models have given significant insight into the structure of real networks. However, real systems almost never reflect the exact details of a model.

Broadly speaking, I see two reasons for the disconnect between theoretical network models, and real-world complex networks. Firstly, networks are fundamentally limited to describing sets of pairwise interactions between nodes – reality is rarely so pairwise. It has become increasingly clear that many systems can only be understood by considering a range of higher-order interactions. This observation has led to rapid growth in the study of higher-order networks including hypergraphs and simplicial complexes [14, 15], although their use is still minimal in comparison to traditional network analysis.

The second reason for the disconnect between the theory and reality of networks is the persis-

tence of a number of key assumptions in the foundations of network theory. Such assumptions include, but are not limited to, treating networks as a tree structure in the absence of loops, assuming the stationarity of certain network properties or rules which determine network evolution, and the framing of network models in terms of node intrinsic properties, excluding node-node interactions. There are excellent reasons for making these assumptions when constructing a theoretical network model, amongst which, in the view of a physicist, is the desire to make network models analytically tractable.

Such a desire is laudable. However, it is my belief that this has resulted in an incomplete understanding of how networks behave if such assumptions do not hold. Exploring these questions is a key motivator for me in pursuing network science.

In the following two chapters, I discuss two key assumptions which are common in network theory. In chapter 2, I consider the role of time dependence in the rules of network growth, specifically, that many approaches to analysing the mechanisms of network growth rely on assuming time independent growth. In chapter 3, I consider node copying models, a mechanism with particular importance to social network formation. In this context, I discuss the assumption of structurally uniform copying probabilities that are common in such models, and define and discuss a model in which structural symmetries in the copying probabilities are broken.

Although the two chapters are largely stand-alone, one idea that is common two both is the concept of deconstructing a complex single layer network model into a two layer graph, which I refer to as a hidden network model.

Formally, a hidden network model can be defined as the pair of single layer graphs $G = (G_O, G_H)$, comprising an observed network $G_O = (V, E_O)$ and a hidden network $G_H = (V, E_H)$, where V is the set of nodes common to both network layers and E_O and E_H are the set of edges for each network. The key feature of a hidden network model is that the evolution of G_O is dependent on G_H (and/or vice versa). Conceptually, hidden network models are closely related to interdependent [16], and multilayer networks [17].

Although there is little about this framework which is novel in mathematical terms, the value

in hidden network models is the conceptual ability to separate our analysis of observed network structures, from the analysis of the underlying (and typically hidden) mechanisms which drive observations. Why is this important? Simply because real networks will often only give an incomplete representation of the real structures which underlie the data (be it a more complex network, or some higher-order structure leading to a lower-order network projection [18]). Hidden network models offer one tool for thinking about such a problem, and allow for a convenient framing of the mathematics underlying each of the models I introduce.

Unfortunately, although I use hidden network models as a tool in the coming chapters, I do not explore their intrinsic properties as much as I would like. Despite this, I urge the reader to consider the role of the hidden/observed layer decomposition as they read on, and how those layers are causally coupled.

It is my belief, although I have not investigated it fully, that the form of the coupling (whether uni- or bidirectional) plays an important role in determining the validity of certain assumptions, for instance, if the rules of network growth can be treated as time independent. I aim to continue pursuing these questions in the future, and hope that the reader will be convinced by the results presented of the value of hidden network models, and that they warrant further investigation.

Chapter 2

Identifying time dependence in network growth

2.1 Publication

Chapter 2 is based on the following publication:

M. Falkenberg *et al.* "Identifying time dependence in network growth". In: *Physical Review Research* 2.2 (2020), p. 023352. DOI: 10.1103/PhysRevResearch.2.023352

Contributions: This work was initiated during my time as a visitor at the Tokyo Institute of Technology in the lab of Yoshihiro Miyake. For this publication I wrote the text and developed the key themes regarding time dependence in network growth. Comments to the manuscript were provided by Tim S Evans, Kim Christensen and Ken-ichiro Ogawa. The code to generate the numerical results was written by me and Jong-Hyeok Lee. The idea to investigate the k2 model was developed by myself, Jong-Hyeok Lee, Shun-ichi Amano and Ken-ichiro Ogawa. Mathematical aspects of the work were developed by myself, Ken-ichiro Ogawa and Tim S Evans. Kim Christensen, Yoshihiro Miyake and Kazuo Yano supervised the project.

2.2 Introduction

Complex networks are typically characterised using a small number of key observables. Amongst these, one of the most popular is the degree distribution [19].

The degree, k, of a node in a network is the number of direct connections a node has to other nodes in the network. Consequently, the degree distribution, $P(k)$, is the probability distribution of the degree across all the nodes in the network, and is said to be scale-free (the exact definition is argued over) if it displays power law scaling such that $P(k) \propto k^{-\kappa}$, where κ is a positive constant, often found to be in the range $\kappa \in \{2, 3\}$ for real networks [20]. Plotting $P(k)$ vs. k on a log-log scale, a power law distribution appears as a straight line with gradient $-\kappa$. In practice, this definition is often loosened such that $P(k) \propto k^{-\kappa}$ for $k \gg 1$ only, accounting for non power-law scaling at small degrees. This distinction is important since the paradigmatic example of scale-free networks, the Barabási-Albert (BA) model, is only scale-free under this loosened definition (the exact scaling is given by the ratio of gamma functions).

Since the late 1990s, many real networks have been reported as having scale-free, or nearly scale-free degree distributions. This includes web-page links on the internet [21], citation networks [22], the co-occurrence of words in language [23], sexual contact [24], social networks [25], and others. Identifying these networks as scale-free has important consequences: (1) it gives a potential mechanistic understanding of the origin and development of these networks, notably that the network evolves according to preferential attachment, and (2) it suggests these networks have a set of important properties associated with scale-free networks. These properties include the presence of hub nodes which have degree much larger than the network average, very small network diameters [26], and resistance to errors but vulnerability to targeted attack [27].

Although the scale-free paradigm has become a hallmark of complex networks research, identifying scale-free behaviour in real networks is still very controversial [28].

Significant effort has gone towards developing appropriate statistical techniques to assess whether networks at a fixed point in time are scale-free, most notably in [19, 29]. Given the difficulty in distinguishing a power law distribution, $P(k) \propto k^{-\kappa}$, from similar distributions such as

the log-normal, $P(k) \propto k^{-1}\exp[-\frac{(\ln k-\mu)^2}{2\sigma^2}]$ (μ: mean degree, σ: standard deviation of degree's natural logarithm), or stretched exponential, $P(k) \propto k^{\tau-1}\mathrm{e}^{-\lambda k^{\tau}}$ (τ: shape parameter, λ: scale parameter), these statistical techniques are clearly important for understanding the statistical properties of network degree distributions.

Applying such techniques to a large set of real world networks, a recent study found that true scale-free networks are rare, representing only about 4% of all networks [19]. These results are broadly in line with a number of similar criticisms of the scale-free paradigm [30–35]. However, despite broad support for these criticisms, many others in the networks community are still strong believers that most complex networks exhibit preferential attachment [36–39]. Among these individuals, many have taken issue with the methods to process the data in [19] and/or the strictness of the scale-free definition [28, 38–40]. Arguing that scale-free networks are only well defined in the infinite system size limit, looser definitions suggest that scale-free networks are in fact not rare at all [39]. However, it can be argued that such a loosening will naturally result in a larger number of positive identifications, and that using weakened criteria for scale-freeness defeats the aim of using a statistically rigorous approach. Clearly, the issue of which approach is best when analysing network degree distributions is yet to be fully resolved.

There is a third camp who argue that "knowledge of whether or not a distribution is heavy-tailed is far more important than whether it can be fit using a power law" [34]. However, great care must be taken with such an approach in a context dependent manner. For instance, in the case of epidemic spreading, two networks may both be fat-tailed with similar degree distributions, yet exhibit very different epidemic mixing patterns due to differences in network assortativity [41, 42].

What all these approaches have in common is that they analyse the degree distribution of a network at a fixed point in time. If such an analysis is to give insight into the mechanistic origin and evolution of a network, it would be prudent to ask whether the degree distribution is representative of the network in general during its evolution, or only for a brief period of time? In other words, is the degree distribution stationary? Without an answer to this question, inferring the past and future evolution of a network based on the current form of its degree

distribution may give misleading results.

A prominent example of a theoretical network model where the observed degree distribution appears to be non-stationary is super-linear preferential attachment, where new nodes attach to existing nodes proportionally to their degree to a power greater than 1 [43]. In the long time limit, a gelation phenomenon is observed where almost all nodes connect to a single hub node forming a star-like network. However, Krapivsky and Krioukov [44] showed that super-linear attachment has significant pre-asymptotic regimes where the degree distribution appears to be approximately scale-free.

Given the difficulty of directly identifying preferential attachment from static degree distributions, proponents of the scale-free paradigm have argued that preferential attachment can be identified directly from dynamical network data (if available) [38]. Numerous approaches have been introduced over the years, using a variety of different assumptions [45–51]. Most commonly, methods assume that the preferential attachment kernel follows a functional form, $\Pi(k) \propto k^{\gamma}$, and primarily focus on estimating the exponent γ - such methods will naturally assume that the preferential attachment kernel of a network is *time independent*.

As an alternative approach, non-parametric methods have been proposed that do not assume a functional form. The first of these methods by Jeong, Néda, and Barabási [46] infers the form of the attachment kernel by constructing a histogram of the degree of nodes to which new edges attach over a short observation window. However, there is no clear guide as to how to choose the start of the observation window and how long it must be - too short and the result is very noisy, too long and the result is subject to bias [51]. The method by Newman [45] avoids this problem by constructing multiple histograms over different observation windows and computes the attachment kernel by taking a weighted average over the different histograms. Although this method avoids the issue of how to choose your observation window, this approach seems to underestimate the attachment kernel at large degrees [52], an issue since corrected by Pham, Sheridan, and Shimodaira [51].

For networks in which the attachment kernel is time independent, the corrected Newman method proposed in [51] gives an excellent fit to data. However, it is still not clearly es-

tablished whether the assumption of time independence is valid for real networks, and in some cases (such as citation networks) it is known to be false [53]. Similarly, the probability of attaching to a node may be a function of a variable other than the degree. However, how to correctly identify which feature of a node determines its attractiveness is not clear.

It is often argued that accurately calculating the attachment kernel of a growing network is important because it can help to predict the future evolution of a network [50]. For instance, in the case of non-linear preferential attachment, where the attachment kernel is given by $\Pi(k) \propto k^{\gamma}$ with positive exponent γ, it is known that for $0 < \gamma < 1$, the limiting degree distribution is a stretched exponential, whereas for $\gamma > 1$, the degree distribution displays a gelation phenomenon where a single dominant hub connects to almost all other nodes in the network [43]. In between, $\gamma = 1$ corresponds to traditional linear preferential attachment where the degree distribution displays power law scaling. Hence, if we can estimate the value of γ for the attachment kernel of a real network, this can be used to predict its future evolution.

Predictions regarding the future evolution of networks, explanations of the historical development of networks, and investigations into whether preferential attachment underlies the evolution of networks, based on measured attachment kernels, are widespread in the literature. These include studies on citation networks [54, 55], protein networks [56], the bitcoin network [57], common words in the English language [58], social dynamics in online games [59], actor networks [46], and more.

The majority of these studies make three assumptions: (1) that the degree, k, of a node is the key feature determining a node's attractiveness, (2) that the attachment kernel can be approximated by $\Pi(k) \propto k^{\gamma}$, and (3) that the measured attachment kernel is either time independent, or that the time dependence is largely unimportant. For instance, looking at four different periods in the evolution of the *American Physical Society* (APS) citation network, and using the node degree (citation count) as the key variable of interest, Sheridan and Onodera found that the exponent γ ranges from 0.94 to 1.06 [55]. The authors assert that this implies that the attachment probabilities in the APS citation network are at least approximately time independent. However, as noted, $\gamma < 1$ would imply that the APS citation network's degree

distribution approaches a stretched exponential, whereas $\gamma > 1$ would result in a gelation effect. Since both $\gamma < 1$ and $\gamma > 1$ were observed from the data, what does this imply for the future evolution of the network?

The issues that arise when making these assumptions can be highlighted more explicitly if we consider a general master equation for a growing network. Let us assume that only one node is added to the network at each time step, and that the probability that the new node connects to an existing node is a function of the existing node's degree and time only. We can now write the following master equation,

$$N(t+1)p(k,t+1) - N(t)p(k,t) = N(t)\Pi(k-1,t)p(k-1,t) - N(t)\Pi(k,t)p(k,t) + \delta_{k,1}, \quad (2.1)$$

where $N(t)$ is the number of nodes at time t, $p(k,t)$ is the probability that a randomly chosen node in the network has degree k at time t, $\Pi(k,t)$ is the probability that at time t an existing node with degree k receives a new edge from the newly added node, and $\delta_{k,1}$ is the Kronecker delta function which represents the addition of a single new node with degree $k = 1$.

Written in this way, the left hand side of Eq. (2.1) represents the change in the number of nodes with degree k from time t to time $t+1$. This is equal to the number of nodes with degree $k-1$ which gain a new edge at time t (first term on right hand side) so that they have degree k at time $t+1$, minus the number of nodes with degree k which gain a new edge at time t (second term) so that they have degree $k+1$ at time $t+1$. The Kronecker delta only applies if $k = 1$ and represents the contribution of the new node.

Without loss of generality, we take $k > 1$ such that $\delta_{k,1} = 0$, and make the following assumption:

$$\text{Let } p(k,t) = p(k,t+1) \text{ and take } t \to \infty \text{ such that } p(k,t) = p_\infty(k).$$

This stationarity assumption is the foundational assumption in the majority of analytically tractable network models. Applied to the Master equation this gives

$$N(t+1)p_\infty(k) - N(t)p_\infty(k) = N(t)\Pi(k-1,t)p_\infty(k-1) - N(t)\Pi(k,t)p_\infty(k). \quad (2.2)$$

In our formulation one node is added to the network at each timestep such that we can take $N(t) = t$ and $N(t+1) - N(t) = 1$, so that Eq. (2.2) can be written as

$$p_\infty(k)\,[1 + t\Pi(k,t)] = p_\infty(k-1)t\Pi(k-1,t), \tag{2.3}$$

which rearranged gives

$$\frac{p_\infty(k)}{p_\infty(k-1)} = \frac{t\Pi(k-1,t)}{1+t\Pi(k,t)}. \tag{2.4}$$

Using this framing, it is easy to see how the stationarity of the degree distribution and the stationarity of the rules of network growth are related: The left hand side of Eq. (2.4) is time independent, whereas, in general, the right hand side is time dependent. Unless, the time dependence on the right hand side cancels out, this result implies that the stationarity assumption is invalid. It is perhaps not a surprise then that the most common analytically tractable growing network models do ensure that the time dependence in Eq. (2.4) cancels.

The aim of this chapter is to illustrate the risks of assuming stationarity in the rules governing the evolution of growing networks, and the risk of assuming that node degree determines node attractiveness. We will do this by introducing the "k2 model", a simple variant of the Barabási-Albert (BA) model where new nodes do not attach to existing nodes proportionally to the number of direct neighbours a node has, but rather proportionally to the number of nodes within a distance two of the target node. This simple rule is rooted in the idea that well connected neighbours are preferable to poorly connected neighbours. The rule puts a particular focus on the role of nearest-neighbour correlations in network growth. Such mechanisms of mutual benefit may be relevant to collaboration [60], or citation [61] networks. The mechanism may also have indirect relevance to node copying processes [62–64]. Similar ideas have been explored in [65–68].

Although this simple rule has no explicit time dependence (i.e., time dependence is not in built by including a time dependent parameter, e.g., the average out-degree), the correlations that form between neighbouring nodes result in an implicit time dependence in the attachment kernel. Consequently, the resulting network does not demonstrate any of the simple scaling

observed in traditional network models. This is despite an extended initial transient phase during which the network appears to grow according to linear preferential attachment. We support this argument with a simple analytical treatment demonstrating that assumptions of simple scaling in the k2 model are not robust.

The arguments we illustrate with the k2 model are highly relevant to real networks. By calculating the ratio of network attachment kernels over different time periods, we show that over short timescales, assumptions of time independence for real networks are relatively well justified. However, over longer time periods, the relative attachment kernels calculated show clear time dependence, displaying a diversity of patterns. While the overall effect may be small in some cases (such as for the Flickr friendship network or the English Wikipedia hyperlink network [69]), we argue that, at a minimum, practitioners should test the degree of time dependence in their data before making predictions about the future or past development of a network.

2.3 Methods

2.3.1 Model definition

The k2 model is defined as a simple, undirected network. The model is initialised with a small connected network of m_0 nodes. Each time-step, a new node is created with $m \leq m_0$ new edges. The m edges are connected to the new node and target nodes from the network. Each target node is chosen with probability proportional to the number of neighbours which are one or two steps away, $k_i^{(2)}$, from the target node, i. We refer to $k_i^{(2)}$ as the second degree of node i. The attachment probability is identical to the BA model with the exception that the BA model attaches proportionally to the number of nodes one step away, $k_i^{(1)}$, from the target node, i. We refer to $k_i^{(1)}$ as the first degree, or just the degree, of node i. Computationally, we prevent multiple edges being formed between two nodes by selecting m unique target nodes.

For clarity, whenever notation is presented with a subscript i or j, for instance $k_i^{(1)}$ or $k_j^{(2)}$,

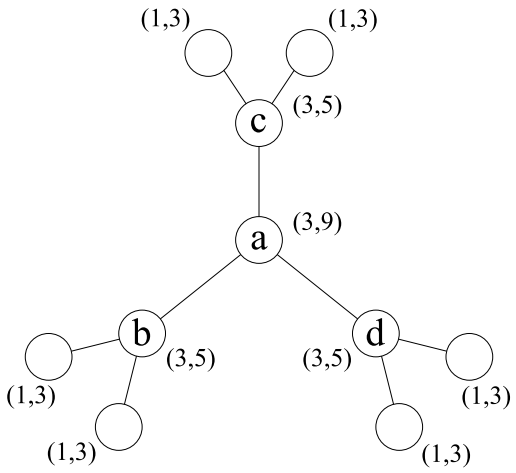

Figure 2.1: A sketch of a simple tree network with four labelled nodes. The values in the brackets correspond to the first and second degree, $(k_i^{(1)}, k_i^{(2)})$, of each node. The four labelled nodes have the same degree, $k_i^{(1)} = 3$, implying equal importance in the BA model, but different second degrees, implying unequal importance in the k2 model.

the focus is on the value of that variable for the particular node i or j. When the subscript is omitted, for instance $k^{(1)}$ or $k^{(2)}$, the focus is on all nodes with the same specific value of the variable in question. We use k and $k^{(1)}$ interchangeably where appropriate.

Figure 2.1 illustrates the motivation for the k2 model. In the BA model, a node's importance is proportional to the number of nodes connected to it, i.e., the first degree. However, there is no consideration for whether these connected nodes are important or not. A node with three isolated neighbours is considered equally important to a node neighbouring three hubs. This is in conflict with many real world scenarios, for instance in academic collaboration networks, where it is known that junior researchers working under top scientists are those most likely to be successful and reach tenure in their careers [60]. In the k2 model, this effect is accounted for, allowing nodes to benefit from connecting to hub nodes and giving them the opportunity to become hubs themselves.

The principle of weighting neighbour importance reflects the role of friends-of-friends in social network theory [70, 71], and is the foundation for widely used network centrality measures built on self-consistent equations, such as Katz centrality [72, 73] or PageRank [73, 74].

Mathematically, we define the attachment kernel, Π, as the function specifying the probability of attaching to a specific node in the network. In the BA model, $\Pi^{(BA)} \propto k^{(1)}$, whereas in the

k2 model, $\Pi^{(k2)} \propto k^{(2)}$. In the case of the k2 model, we can write the normalised form of the attachment kernel as

$$\Pi_i^{(k2)} = \frac{k_i^{(2)}}{\sum_{j=1}^{N} k_j^{(2)}} \approx \frac{k_i^{(2)}}{\sum_{j=1}^{N} (k_j^{(1)})^2}, \tag{2.5}$$

where for $m = 1$, the approximation is an equality. By splitting the numerator of the attachment kernel into the contribution of the first degree neighbours to node i, $k_i^{(1)}$, and the contribution of the next-nearest-neighbours, $k_i^{(2)} - k_i^{(1)}$, Eq. (2.5) can be rewritten as

$$\Pi_i^{(k2)} \approx \frac{k_i^{(1)} + \sum_{\alpha=1}^{k_i^{(1)}} (k_{i\alpha}^{(1)} - 1)}{\sum_{j=1}^{N} (k_j^{(1)})^2}, \tag{2.6}$$

which is a function of first neighbour degrees only, where we have used

$$k_i^{(2)} = \sum_{\alpha=1}^{k_i^{(1)}} k_{i\alpha}^{(1)}. \tag{2.7}$$

Here, α labels the $k_i^{(1)}$ unique first neighbours of node i, and $k_{i\alpha}^{(1)}$ is the first degree of node α, connected to node i. In Eq. (2.6), the first term indicates the contribution to the attachment kernel from the direct neighbours of node i, and the second term indicates the contribution from the next-nearest neighbours of node i.

Conceptually, we can think of the k2 model as involving two separate network layers forming what I refer to as a *hidden network model*. In the *observed* network, each node represents an agent, and an edge between two nodes represents a direct, first degree relationship between the two nodes. However, new nodes do not connect to a target node according to the node's direct connections, but rather according to the number of nodes within distance two of the target. These nodes are within the sphere of influence of the target node. Hence, we define the *influence* network (which I later refer to as the *hidden* network,[1]) in which an edge between any two nodes signifies that the nodes are within each other's sphere of influence, i.e., two connected nodes are neighbours, or next-nearest-neighbours in the observed network, see Fig. 2.2.

[1]Note, I have retained the "influence network" terminology in this chapter to ensure consistency with the published version in Ref. [1].

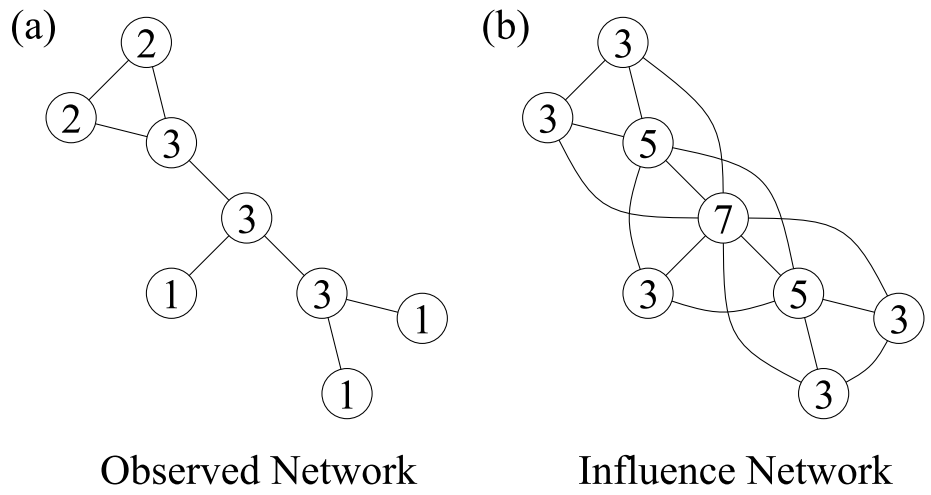

Figure 2.2: A schematic showing how the k2 model can be decomposed into two distinct network layers. (a) The *observed* network is the network of nodes which have a direct connection to each other. In this network, the degree does not account for the importance of next-nearest-neighbours. (b) The *influence* network where two nodes are connected if they are nearest or next-nearest-neighbours in the observed network.

The influence network has similarities to the node copying mechanism studied in [62, 63], based on earlier models in [64], with relevance to social network formation [70, 75], citation networks [61, 76], evolution [77], and protein interaction networks [78, 79]. Although the k2 model is not designed to model such systems explicitly, it may be useful for understanding the role of neighbour-neighbour correlations in the growth of such networks. A variant of this idea is studied in chapter 3 to investigate heterogeneous node copying.

In the influence network, new nodes connect to a target node proportionally to the node's degree, $k^{(2)}$. The new node then copies a fraction of the nodes attached to the initial target node, and forms additional edges to these copied neighbours. The copied neighbours correspond to those which are directly connected to the target node in the observed network. This illustrates a key property of the k2 model – when decomposed into a hidden network model, the influence/hidden network determines the future evolution of the observed network via direct attachment (hidden causally affects observed), whereas the observed network determines the future copying dynamics in the influence/hidden network (observed causally affects hidden), indicating the presence of bidirectional coupling between the two layers of the hidden network model decomposition.

2.3.2 Measuring the time dependence of preferential attachment

To understand how the attachment kernel of a network changes over time, it is helpful to consider relative attachment probabilities as opposed to absolute attachment probabilities. In general we can write an arbitrary attachment kernel, which is a function of the node degree only, as

$$\Pi(k;t) = \frac{f(k)}{\sum_{j=1}^{N(t)} f(k_j(t))} \tag{2.8}$$

with an arbitrary preference function f. The summation is over all nodes in the network, $N(t)$, at time t. The function f is *time independent*, however, as the network grows and more nodes are added, $N(t)$ in the denominator changes, and hence, the denominator is *time dependent*. Note, $f(k_i(t))$ for a specific node i is *time dependent*, since the degree of a specific node evolves over time. We define the relative attachment kernel as

$$\phi_t(k,k') = \frac{\Pi(k;t)}{\Pi(k';t)} = \frac{f(k)}{f(k')}. \tag{2.9}$$

As opposed to $\Pi(k;t)$, the relative attachment kernel has no dependence on the network as a whole, but rather, is a function of the degree k and k' only. As a result, we can express the time independence of the attachment kernel as

$$\frac{\mathrm{d}\phi_t(k,k')}{\mathrm{d}t} = 0. \tag{2.10}$$

For convenience, in the following we will consider $\phi_t(k,1)$, i.e., the attachment probability of connecting to a node with degree k relative to a node with degree $k' = 1$. By definition, $\phi_t(k,k) = 1$.

Consider the relative attachment kernel calculated at time t, written as $\phi_t(k,1)$, and at time s, $\phi_s(k,1)$. If Eq. (2.10) holds, then $\phi_t(k,1) = \phi_s(k,1)$. For a real network, it is likely that there will be small deviations from this ideal case. Hence, we can plot the ratio $\phi_t(k,1)/\phi_s(k,1)$ against degree k to gauge the extent of the time dependence across a specific time interval. This ratio is only well defined for networks which contain nodes with degree k at both times t

and s.

The BA model is a simple case where Eq. (2.10) should hold, with $\phi_t(k,1) = k$ for all t. Likewise, for non-linear preferential attachment, $\phi_t(k,1) \propto k^\gamma$ with positive constant γ. In the case of the k2 model, Eq. (2.10) does not hold, due to the second term in the numerator in Eq. (2.6). The preference function in the k2 model is not a function of the degree of a node, but the second degree, $k^{(2)}$. Clearly the second degree is related to the first degree, and, when analysing the k2 model, one could mistakenly believe that the node degree, $k^{(1)}$ is the quantity determining network growth. However, although this appears approximately true at first, over time, the relation between the first and second degree changes; $\left\langle k^{(2)}(t) \right\rangle$ is not proportional to $\left\langle k^{(1)}(t) \right\rangle$. In other words, although the attachment kernel is not explicitly time dependent (e.g., we have not included an explicit ageing mechanism), the local network structure, which determines a node's second degree, is time dependent.

This point cannot be overstated; while the k2 model clearly breaks the assumptions outlined above, it does so in a way that, without prior knowledge of the model rules, is wholly non-obvious. As we will outline, if sufficient care is not taken, these assumptions risk misleading or incorrect predictions about a networks past or future evolution.

2.4 Results

2.4.1 Simulation results

In the following, we will focus on analysing the attachment kernel and the degree distribution of the k2 model, using the BA model as a comparison. Each simulation is initialised with a complete graph of $m_0 = m + 1$ nodes.

Figure 2.3 shows the degree distribution and the true relative attachment kernel for the k2 model with $m = 1$. Both sub-figures are averaged over 100 simulations. Early in the network development, there are only small differences between the behaviour of the k2 and BA models. However, as the network grows, significant differences emerge. The duration of the initial

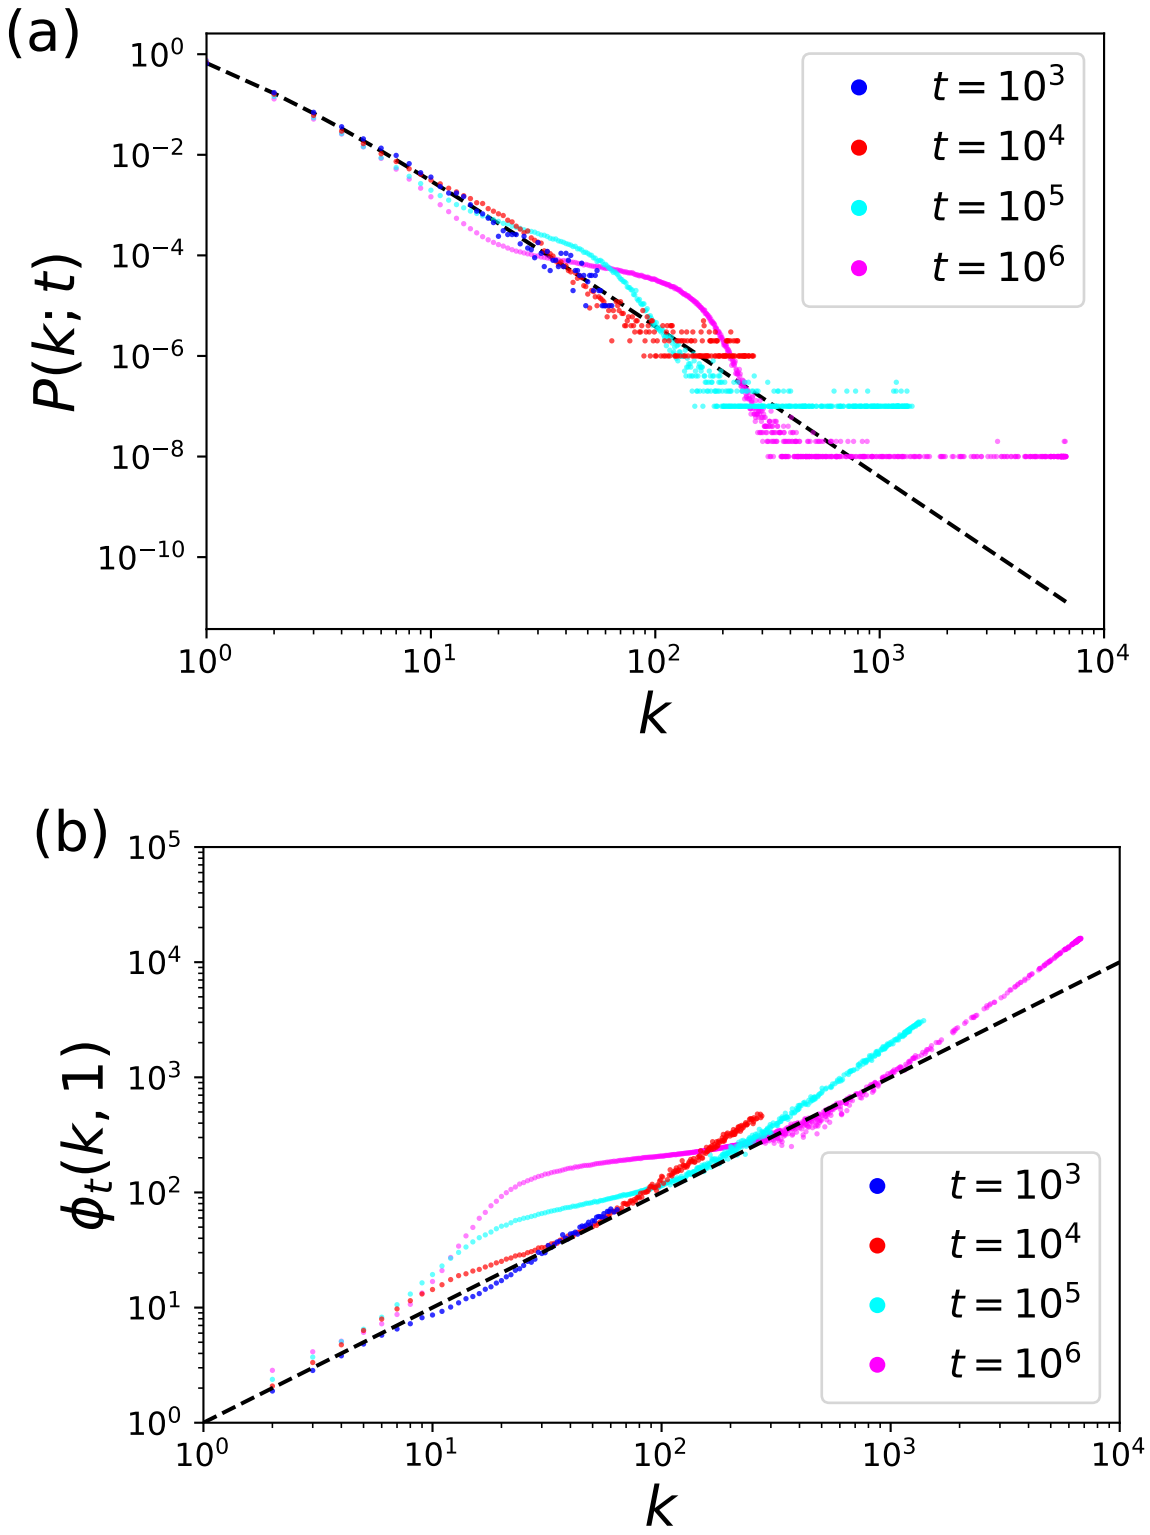

Figure 2.3: The (a) degree distribution and (b) relative attachment kernel for the k2 model with $m = 1$. The dashed lines show the expected scaling for the BA model. Early in the growth of the k2 model, the evolution of the network is largely indistinguishable from the BA model. As the network grows, both the degree distribution and relative attachment kernel deviate significantly from the simple BA model scaling. Data are averaged over 100 simulations. Error bars omitted for clarity.

BA-like transient phase is longer and follows the BA model even more closely for $m > 1$, see [1].

Over short timescales, the network growth appears largely indistinguishable from linear preferential attachment. However, over longer timescales, the attachment kernel shows clear deviations from this simple scaling, with a plateau region at moderate degree preceding a super-linear tail. The anomalous scaling observed is most clearly seen for nodes with moderate degree in the range $k \approx 10$ to $k \approx 300$. This region suggests that there may be multiple timescales of interest at play in the evolution of the k2 model.

As noted, the BA model incorporates a rich-get-richer mechanism, but does not account for any mechanisms of mutual benefit between nodes; new nodes added to the network receive no benefit from attaching to a hub node as opposed to any other less important node. Conversely, in the k2 model, when a node i, added to the network at time t_i, attaches to a hub node, the new node's initial attractiveness is given by

$$\Pi_i(k_i; t_i) \propto k_i^{(2)}(t_i) = 1 + k_{\text{hub}}^{(1)}(t_i) \approx k_{hub}^{(1)}(t_i), \tag{2.11}$$

which is completely determined by the first degree of the targeted hub. This has the counter-intuitive effect that the tail of the attachment kernel appears to show super-linear preferential attachment, implying gelation, but that the change in the number of new edges in the influence network is dominated by new nodes with degree $k^{(1)} = 1$. As a consequence, the k2 model appears to show a gelation-like phenomenon to communities rather than hubs, resulting in the plateaus shown in Fig. 2.3.

The true relative attachment kernel in Fig. 2.3 is typically not accessible for a real network. To illustrate the risks this may pose, let us assume the k2 model is a real network and fit the relative attachment kernel, on the assumption that $\phi_t(k, 1) \propto k^\gamma$ for a positive exponent γ. From Fig. 2.3(b) we may deduce that for $t = 10^3$, the k2 model has an approximately linear (or possibly slightly sub-linear) attachment kernel, whereas at $t = 10^6$, the attachment kernel is highly non-linear, but clearly grows faster with k than the simple prediction from

linear preferential attachment. If we were to use these results to infer the future scaling of the network, the data at $t = 10^3$ would suggest the network might approach a stretched exponential degree distribution, whereas from the data for $t = 10^6$, we might paradoxically infer the network is approaching a gelation state. In the case of the k2 model this approach is misleading, but for other networks this approach may be a good first approximation. However, what is clear is that simply calculating the attachment kernel of a network at one point in time is not sufficient to determine the form of the attachment kernel in the past or future. Likewise, since the degree distribution is determined by the underlying dynamical process growing the network, we cannot accurately know how the degree distribution will evolve in the future.

In the case of a real network, we can only estimate the relative attachment kernel by observing the degree of nodes to which new nodes added to the network attach. To simulate this real-network scenario, we apply the corrected Newman method to a single simulation of the k2 model as shown in Fig. 2.4.[2]

Figure 2.4(a) shows the calculated relative attachment kernel for the k2 model at times $t = 10^5$ and $s = 10^6$. As in Fig. 2.3(b), nodes with moderate degree, $k \approx 30$, show an excess in the relative attachment kernel. In Fig. 2.4(b), deviations in the relative attachment kernel are shown explicitly by taking the ratio to the relative attachment kernels at $t = 10^5$ and $s = 10^6$. For very small degree nodes, the ratio is approximately one indicating that the attachment kernel is approximately time independent at these degrees. Above $k = 10$, the ratio clearly deviates from one, indicating that the relative attachment kernel is time dependent. For visual clarity, Fig. 2.4(c) & (d) are equivalent to (a) & (b), but using the cumulative sum of the relative attachment kernel, defined as

$$\tilde{\phi}(k, 1) = \sum_{\tilde{k} \leq k} \phi(\tilde{k}, 1). \tag{2.12}$$

It is important to note that the estimated attachment kernel using Newman's method is not fully consistent with the true attachment kernel; the magnitude of the excess in the attachment

[2]Note, the corrected Newman method does not output error bars for the calculated attachment kernel.

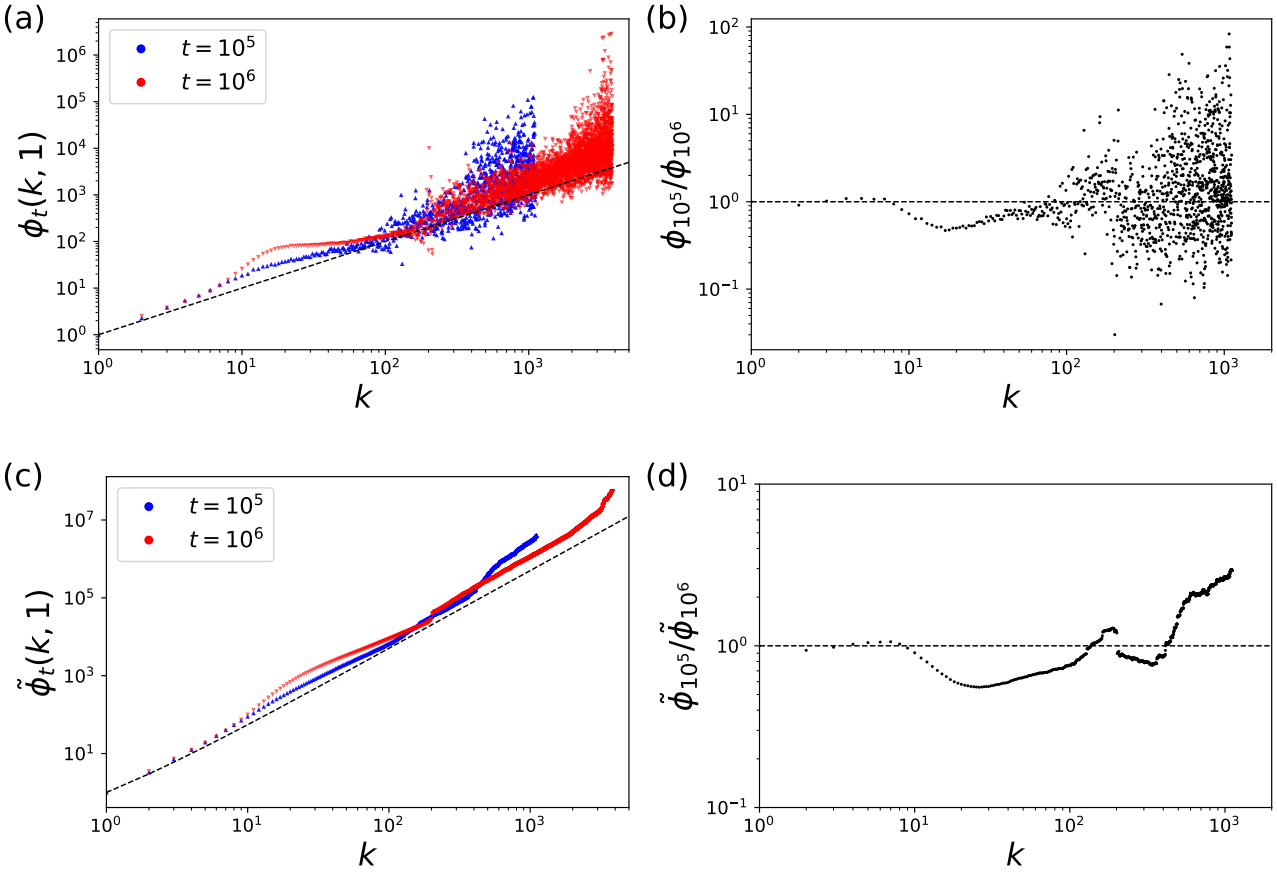

Figure 2.4: An illustration that the relative attachment kernel in the k2 model is *time dependent*. (a) The relative attachment kernel, $\phi_t(k, 1)$, for the k2 model with $m = 1$ as calculated using the corrected Newman method. (b) The ratio of the calculated relative attachment kernels at time $t = 10^5$ and $s = 10^6$. (c) The cumulative sum and (d) corresponding ratio of the relative attachment kernels shown in (a). The dashed lines indicate the prediction for the BA model in (a) and (c), and the ratio expected for any time independent relative attachment kernel, which is a function of node degree only, in (b) and (d).

probabilities is much smaller using Newman's method than the true excess shown in Fig. 2.3(b). This is because Newman's method constructs the attachment kernel by collating multiple histograms from different times in the network evolution. The consequence is that the estimated form of the attachment kernel at $s = 10^6$ is more consistent with the true attachment kernel earlier in the evolution of the k2 model than with the current value of the attachment kernel.

To verify that the deviations in the relative attachment kernel are due to the evolution of the k2 model and not numerical errors, we repeat the analysis shown in Fig. 2.4 for the BA model where the relative attachment kernel is expected to be time independent. Figure 2.5(a) shows that the relative attachment kernels are effectively indistinguishable at different times in the

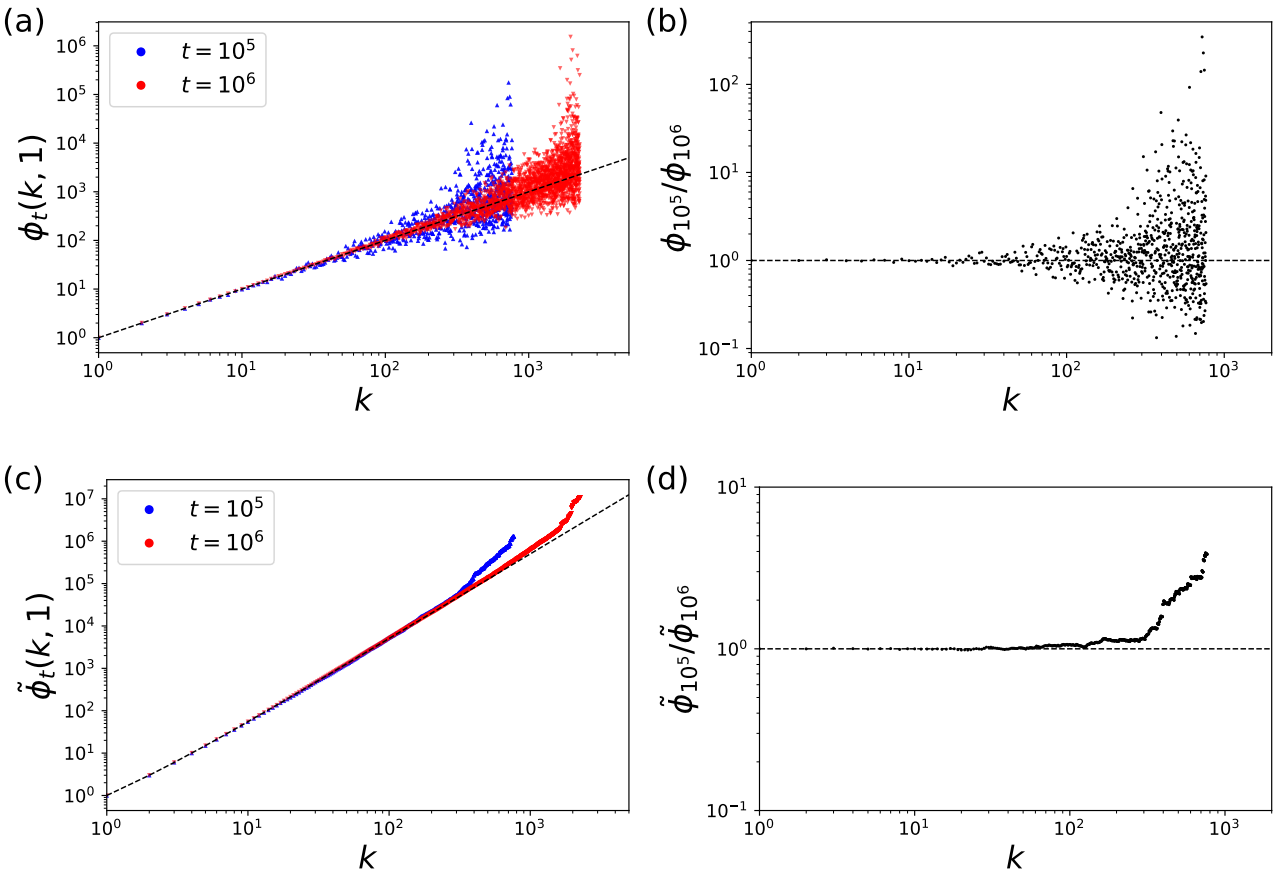

Figure 2.5: An illustration that the relative attachment kernel in the BA model is *time independent*. (a) The relative attachment kernel, $\phi_t(k, 1)$, for the BA model with $m = 1$ as calculated using the corrected Newman method. (b) The ratio of the calculated relative attachment kernels. (c) The cumulative sum and (d) corresponding ratio of the relative attachment kernels shown in (a). The dashed lines indicate the prediction for the BA model in (a) and (c), and the ratio expected for any time independent relative attachment kernel, which is a function of node degree only, in (b) and (d). Finite-size effects result in small deviations from the expected scaling at large $k > 100$, where the data is noisy.

network evolution. This is confirmed by Fig. 2.5(b) where the ratio of the relative attachment kernels is approximately one for all nodes with degree $k < 100$. Noise in the tail obscures the ratio for $k > 100$.

Overall, Fig. 2.5 indicates that using the corrected Newman method is, to an extent, effective at estimating the relative attachment kernel of a network and testing whether it exhibits time dependence. This suggests that the deviations in the relative attachment kernel observed in Fig. 2.4 are due to the structural properties of the k2 model, and not due to limitations in the method used to estimate the relative attachment kernel. Hence, we can deduce that the relative attachment kernel for the k2 model is *time dependent*.

2.4.2 Mathematical results

Given the complexity of the k2 model (specifically the bidirectional coupling between the observed and influence/hidden networks), exact analytical results are hard to derive. However, using simple arguments, we can demonstrate the inconsistencies that arise from assuming the k2 model follows a simple form of non-linear preferential attachment.

From the definition of the k2 model, we can make a continuum approximation and write the evolution of the degree, $k_i^{(1)}(t)$, of node i as

$$\frac{dk_i^{(1)}(\mathrm{t})}{dt} = m\Pi_i^{(k2)}(t), \tag{2.13a}$$

$$\Pi_i^{(k2)}(t) \equiv \frac{k_i^{(2)}(t)}{\sum_{j=1}^{N} k_j^{(2)}(t)} \approx \frac{k_i^{(2)}(t)}{\sum_{j=1}^{N} (k_j^{(1)}(t))^2}, \tag{2.13b}$$

where node i is added to the network at time $t_i \leq t$. The second degree, $k_i^{(2)}(t)$, is defined according to Eq. (2.7), the summation is over all nodes in the network, and the approximation is an equality if $m = 1$; if $m > 1$ the network is no longer a tree and loops result in the double counting of nodes. We can write the evolution of the second degree as

$$\frac{dk_i^{(2)}(\mathrm{t})}{dt} \approx m\frac{\left(k_i^{(1)}(t)\right)^2 + \xi_i(t)}{\sum_{j=1}^{N}\left(k_j^{(1)}(t)\right)^2}, \tag{2.14}$$

where

$$\xi_i(t) \equiv \sum_{\beta=1}^{k_{i_1}^{(1)}(t)-1} k_{i_{1(\beta)}}^{(1)}(t) + \cdots + \sum_{\beta=1}^{k_{i_\alpha}^{(1)}(t)-1} k_{i_{\alpha(\beta)}}^{(1)}(t), \tag{2.15}$$

see section 2.8 for a derivation. Here $i_{\alpha(\beta)}$ represents the node β connected to node i_α, and i_α represents the direct neighbours α connected to node i. Equation (2.15) represents the effect of non-neighbouring nodes on node i. Since the first degree of a node can only grow over time, $\xi_i(t)$ is a positive semi-definite monotone increasing function with respect to time t, that is,

$$\xi_i(t) \geq 0, \tag{2.16a}$$

$$\frac{d}{dt}\xi_i(t) \geq 0. \tag{2.16b}$$

Our aim in the following is to write $\xi_i(t)$ as a function of the first degree, $k^{(1)}$, only. To do so we rearrange Eq. (2.14) to make $\xi_i(t)$ the subject and substitute in Eq. (2.13a) and Eq. (2.13b),

$$\xi_i(t) = \frac{1}{m}\left(\sum_{j=1}^{N}\left(k_j^{(1)}(t)\right)^2\right)\cdot\frac{d}{dt}\left[\frac{1}{m}\left(\sum_{j=1}^{N}\left(k_j^{(1)}(t)\right)^2\right)\cdot\frac{dk_i^{(1)}(\mathrm{t})}{dt}\right] - \left(k_i^{(1)}(t)\right)^2, \tag{2.17}$$

which is a function of the first degree only. Here we note that the summations in Eq. (2.17) correspond to the denominator in Eq. (2.13b). This is the sum over $k_i^{(2)}(t)$ for each node in the network, and hence, corresponds to the twice the total number of edges in the influence network, which we label as

$$E^{(2)}(t) = \frac{1}{2}\sum_{j=1}^{N(t)} k_j^{(2)}(t) \approx \frac{1}{2}\sum_{j=1}^{N(t)}(k_j^{(1)}(t))^2. \tag{2.18}$$

Figure 2.3 shows the degree distribution and the relative attachment kernel for the k2 model obtained from simulations. As a thought experiment, let us suppose that these simulations are not for a theoretical network model but that the data represents a real world network. For the network at small times in its evolution, the degree distribution and the attachment kernel are closely approximated by the BA model.

For preferential attachment where the attachment kernel is a function of degree only, in the absence a constant offset (i.e., where $\Pi \propto k^\gamma$, not $\Pi \propto k^\gamma + a$, with constant a), it is known that, on average, the degree of a given node i evolves in time as a power function given by

$$k_i^{(1)}(t) = m\left(\frac{t}{t_i}\right)^\delta,\ t \geq t_i, \tag{2.19}$$

where t_i is the time at which node i was added to the network, and $\delta = 1/2$ for linear preferential attachment [73]. In the case of sub- (super-) linear preferential attachment, $\delta < 1/2$ ($\delta > 1/2$). Note similar results are found for the fitness model ($\Pi \propto k + a$, with constant a), without the need for a sub- or superlinear attachment kernel. However, assuming an attachment kernel in the absence of a constant offset is common in the literature estimating real-world attachment kernels (see [50]), and is sufficient for illustrative purposes.

Let us assume Eq. (2.19) holds and test whether this simple scaling is consistent with the mathematical form of the k2 model. First, we substitute Eq. (2.18) into Eq. (2.17),

$$\xi_i(t) = \frac{4}{m^2} E^{(2)}(t) \frac{d}{dt} \left[E^{(2)}(t) \frac{dk_i^{(1)}(t)}{dt} \right] - (k_i^{(1)}(t))^2. \tag{2.20}$$

In the case of the k2 model where one node is added to the network at each time step, we initialise our network such that $t_j = j$ and note that the number of nodes in the network at time t is given by $N(t) = m_0 + t \approx t$ for large t.

Using this initialisation, we now calculate the value of $E^{(2)}(t)$ by approximating the sum as an integral and substituting in Eq. (2.19),

$$\begin{aligned} E^{(2)}(t) &\approx \frac{1}{2} \sum_{j=1}^{t} (k_j^{(1)}(t))^2 \\ &\approx \frac{m^2}{2} \int_1^t dt' \left(\frac{t}{t'} \right)^{2\delta} \\ &\approx \frac{m^2 t^{2\delta}}{2(1-2\delta)} \left[(t')^{1-2\delta} \right]_1^t . \end{aligned} \tag{2.21}$$

There are three cases for the different possible values of δ: Case (i) $2\delta < 1$, corresponding to sub-linear preferential attachment, this scenario is likely to be irrelevant for the k2 model since the influence network cannot grow slower than the original network in the BA model. In this case we expect to find linear growth in the number of edges in the influence network

$$E^{(2)}(t) \approx \frac{m^2}{2(1-2\delta)} t. \tag{2.22}$$

Here $E^{(2)}(t)$ is dominated by the youngest nodes (created at the largest times t_i) as the older nodes grow too slowly.

Case (ii) $2\delta = 1$, corresponding to linear preferential attachment,

$$E^{(2)}(t) \approx \frac{m^2}{2} t \ln(t). \tag{2.23}$$

Case (iii) $2\delta > 1$, corresponding to super-linear preferential attachment where there is some enhancement over linear preferential attachment. For the k2 model, this scenario is plausible since we know that for any given node $k_i^{(2)}(t) \geq k_i^{(1)}(t)$. In this case we find

$$E^{(2)}(t) \approx \frac{m^2}{2(2\delta - 1)}(t^{2\delta} - t), \tag{2.24}$$

where the growth in the number of edges in the influence network is dominated by the oldest nodes in the network.

Let us assume case (iii) is valid for the k2 model. Substituting Eq. (2.24) into Eq. (2.20) we find,

$$\begin{aligned} \xi_i(t) &= m^3 \frac{\delta}{(2\delta-1)^2}\left(\frac{1}{t_i}\right)^{\delta}\left[(3\delta-1)t^{5\delta-2} - (4\delta-1)t^{3\delta-1} + \delta t^{\delta}\right] - m^2\left(\frac{t}{t_i}\right)^{2\delta} \\ &= a_1 t^{5\delta-2} + a_2 t^{\delta} - a_3 t^{3\delta-1} - a_4 t^{2\delta}, \end{aligned} \tag{2.25}$$

where in the final line we have grouped all the constants for each term into a single positive prefactor, a_1 to a_4.

Recall that the k2 model requires that $\xi_i(t)$ is a positive, semi-defined monotonically increasing function, and note that Eq. (2.25) is only valid for $\delta > 1/2$. As $t \to \infty$, the first term of Eq. (2.25) will dominate the second if $5\delta - 2 \geq \delta$, $\delta \geq 1/2$. Likewise, the third term will dominate the fourth if $3\delta - 1 \geq 2\delta$, $\delta \geq 1$. Hence, as $t \to \infty$, the first term is the dominant positive term and the fourth term is the dominant negative term. To ensure $\xi_i(t) \geq 0$ for all $t > t_i$, we require that the first term grows faster than the fourth term giving $5\delta - 2 \geq 2\delta$, corresponding to $\delta \geq 2/3$.

Returning to Eq. (2.14) and substituting in Eq. (2.19) and Eq. (2.24), we can also write

$$\frac{dk_i^{(2)}(\mathrm{t})}{dt} = m(2\delta - 1)\frac{m^2\left(t/t_i\right)^{2\delta} + \xi_i(t)}{m^2(t^{2\delta} - t)}. \tag{2.26}$$

We have established that to satisfy Eq. (2.16a), the leading term of $\xi_i(t)$ must scale as $t^{5\delta-2}$, and $\delta \geq 2/3$. However, as a consequence of the rules of the k2 model, at time $t > t_i$, node

i can gain no more than m new edges in the influence network in any given time step (i.e., $k_i^{(2)}(t+1) - k_i^{(2)}(t) \leq m$). Hence, strictly for $t > t_i$, we require

$$\frac{dk_i^{(2)}(\mathrm{t})}{dt} \leq m, \tag{2.27}$$

which is only satisfied if the denominator of Eq. (2.26) grows at least as fast as the numerator of Eq. (2.26). This requires $t^{2\delta} \geq t^{5\delta-2}$ as $t \to \infty$. Hence, $2\delta \geq 5\delta - 2$, giving $\delta \leq 2/3$. Combining the conditions in Eq. (2.16a) and Eq. (2.27), we find that a power function of the form given in Eq. (2.19) can only satisfy the requirements of the k2 model if $\delta = 2/3$.

To test the validity of our argument, we simulate the growth in the number of edges for the k2 influence network. This is shown in Fig. 2.6 for $m = 1$ & 3. The figure shows that, at large t, the number of edges in the influence network scales as approximately $t^{4/3}$ corresponding to $\delta = 2/3$, in agreement with our prediction.

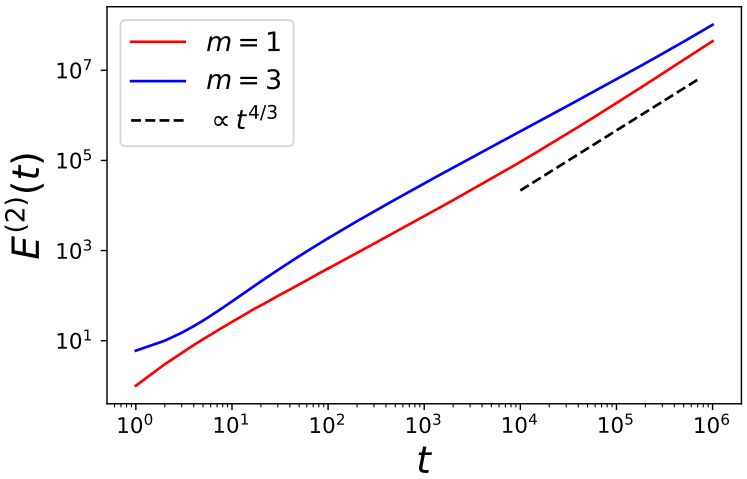

Figure 2.6: The number of edges in the k2 influence network. Data is averaged over 100 simulations. Error bars are negligible. For large t, the number of edges scales approximately with $t^{4/3}$, as indicated by the dashed line.

However, further analysis appears to contradict this conclusion. Firstly, we can simulate the k2 model and track the degree of specific nodes over time, see Fig. 2.7. The data has been averaged over 10^4 simulations with the shaded regions indicating the standard deviation; only with a very large sample size can the average evolution of node i be observed. In most simulations, a node hardly grows at all, whereas in a few simulations, nodes grow very quickly.

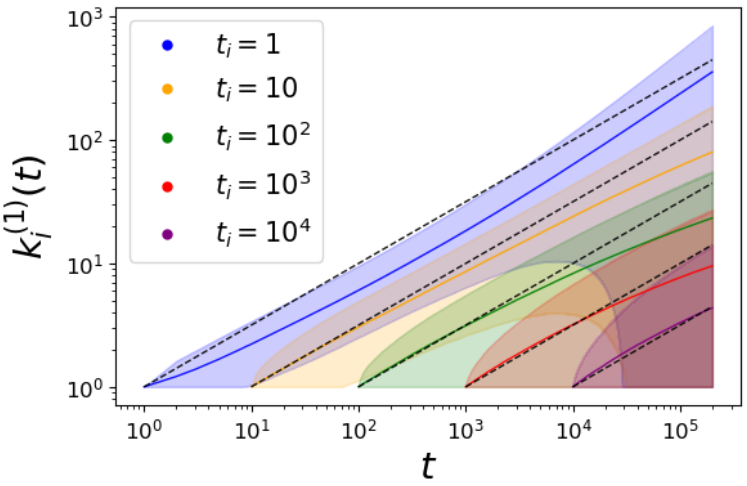

Figure 2.7: The evolution of the degree of individual nodes added at time t_i for the k2 model with $m = 1$. The shaded region around the solid lines indicates the standard deviation across 10000 simulations. The dashed lines indicate the scaling expected for the BA model, see Eq. (2.19).

For a transient period after being added to the network, the average degree evolution of a node appears to scale as $t^{1/2}$ which is the expected scaling for linear preferential attachment. This appears to contradict the $\delta = 2/3$ scaling identified previously, although we note that the integral in Eq. (2.21) is dominated by the oldest nodes in the network for $\delta > 1/2$, which do appear to grow faster than $t^{1/2}$ towards the end of the simulation.

For newer nodes, after a transient period, the degree evolution appears to deviate from $t^{1/2}$ scaling, but the scaling appears to transition to $\delta < 1/2$ rather than $\delta > 1/2$. The time over which this transition takes place increases with the time nodes are added to the network.

This suggests that the true functional form for the degree evolution in the k2 model involves two competing terms, the first scaling as $t^{2/3}$ which is suppressed by t_i, and a second term which scales as $t^{1/2}$ which is suppressed by t. We hypothesise, but at this stage cannot prove, that this implies two scaling regimes: For fixed t_i and $t \to \infty$, the scaling of the degree evolution is dominated by a $t^{2/3}$ term to ensure that $E(t) \propto t^{4/3}$ as $t \to \infty$. For $t_i \to \infty$ and $t = t_i + \epsilon$ where $\epsilon \ll t_i$, the degree evolution of node i is dominated by a $t^{1/2}$ term.

It is interesting to consider the origin of the $t^{1/2}$ scaling. Our results are inconclusive, however,

if we let $t_i \to \infty$ and set $t = t_i + \epsilon$ with $\epsilon \ll t_i$, a Taylor expansion of Eq. (2.25) gives

$$\xi_i(t_i) = b_1 t_i^{4\delta-2} - b_3 t_i^{2\delta-1} + \mathcal{O}(\epsilon), \tag{2.28}$$

with positive constants b_1 and b_3, revealing that $\delta \geq 1/2$, rather than $\delta \geq 2/3$, is sufficient for ensuring that $\xi_i(t = t_i) \geq 0$ as $t_i \to \infty$.

The mathematical argument presented here does not prove the limiting behaviour of the k2 model. However, the result does indicate that the inclusion of simple nearest-neighbour correlations in network growth can affect the scaling of key observables. Despite initially appearing to grow as linear preferential attachment, this simple scaling breaks down as the network grows. In the case of real networks this may happen at an early stage in the evolution of a network. However, as illustrated by the k2 model, the transient time during which the model appears to grow according to linear preferential attachment may be significant – it is not uncommon to analyse real networks with $10^4 - 10^5$ nodes, yet in the case of the k2 model, the network is still in this transient period. In cases like the k2 model with complex growth rules, oversimplified assumptions derived from non-linear preferential attachment do not reflect reality.

2.4.3 Application to real world networks

Figure 2.8 shows the ratio of the cumulative relative attachment kernels for six real world growing networks: (a) a regional friendship network on Facebook [80], (b) the Youtube follower network [69], (c) the APS citation network [81], (d) the `hep-ph` arXiv collaboration network [50], (e) the Flickr follower network [69], and (f) the hyperlink network for English Wikipedia [69]. For all networks, the cumulative relative attachment kernel is calculated at two time points early and late in the network's evolution relative to the endpoint of the dataset. In some cases the full evolution of the network is not known and is accounted for with a large initial graph at $t = 0$.

Two details are clear in Fig. 2.8. Firstly, on short timescales the relative attachment kernels are approximately time independent, with only small deviations observed. However, over

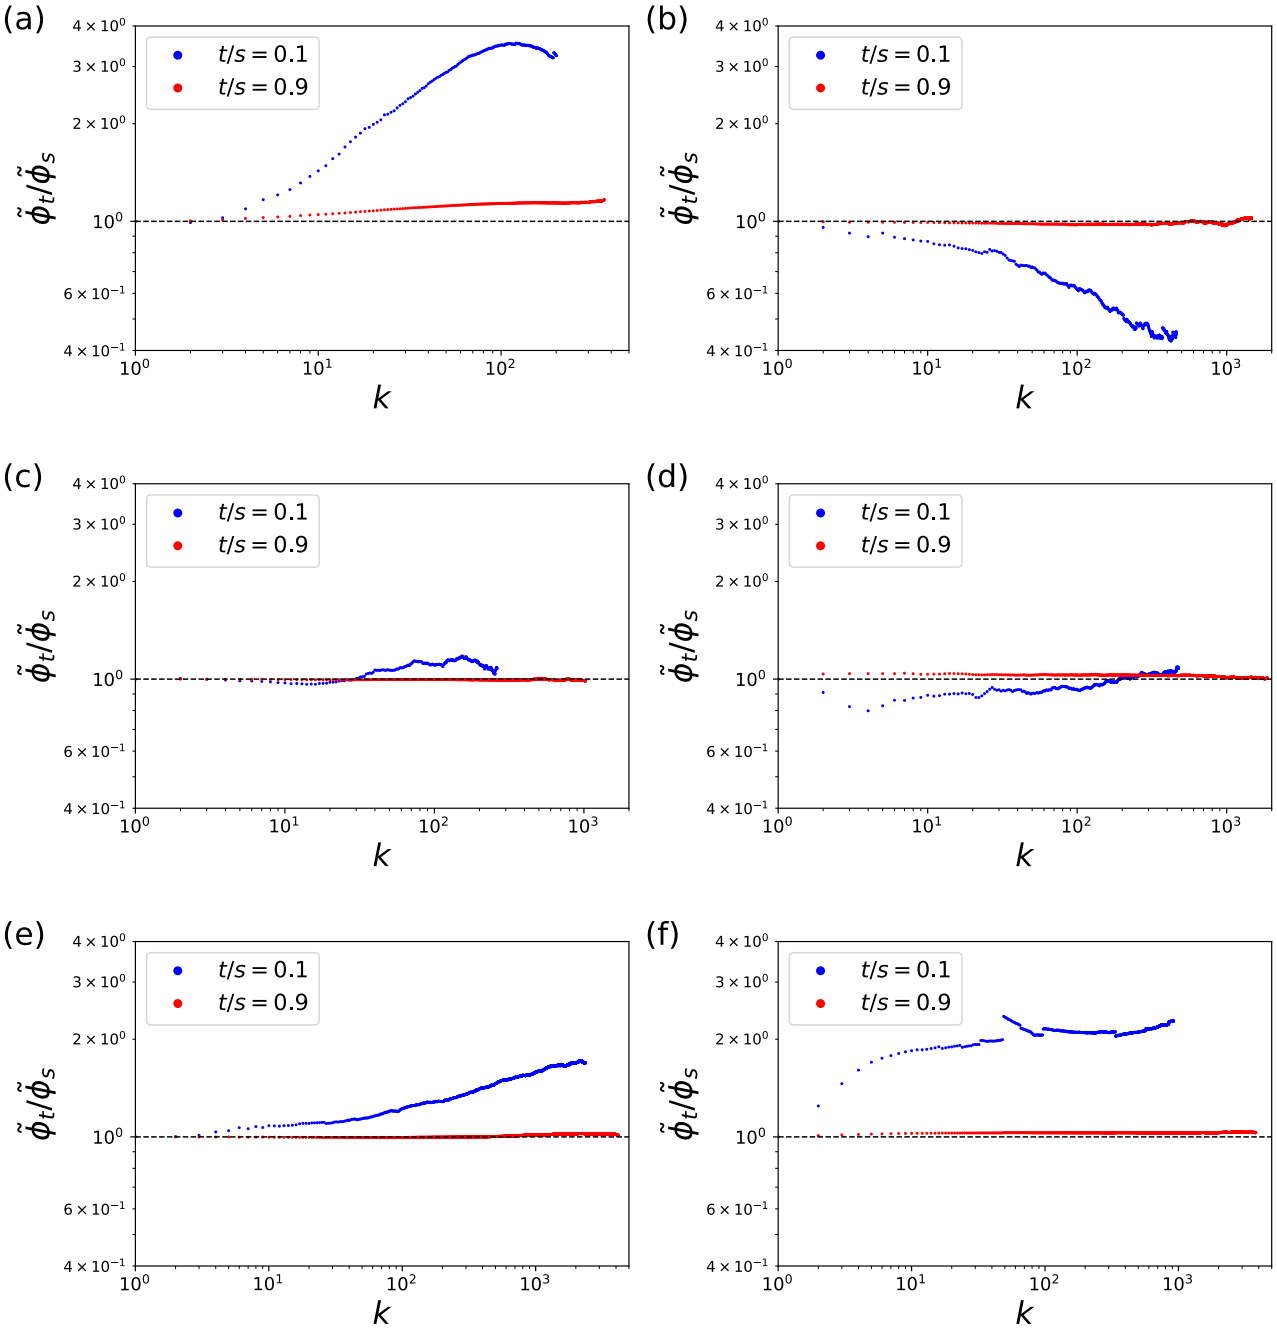

Figure 2.8: The ratio of the cumulative relative attachment kernels for six real world networks; (a) Facebook friendships [80], (b) Youtube followers [69], the APS citation network [81], (d) `hep-ph` arXiv collaborations [82], (e) the Flickr network [69], and (f) hyperlinks on English Wikipedia [69]. For all networks the ratio is shown for early, $t/s = 0.1$, and late, $t/s = 0.9$, in the evolution of the recorded network relative to the endpoint, s; time is measured in the net number of edges added to the network.

longer time periods, the ratio of the relative attachment kernels is not constant indicating time dependence.

There is significant diversity in the changes observed to the ratio of the relative attachment kernels. In Fig. 2.8(a), (e) and (f), the ratio is (to a good approximation) monotonically increasing. This implies that if we were to approximate the attachment kernel of these networks using non-linear preferential attachment, $\Pi(k) \propto k^{\gamma}$, the exponent γ will appear to have reduced over time (the network growth is becoming more sub-linear). Conversely, in Fig. 2.8(b) we see the opposite effect where the ratio is approximately monotonically decreasing, implying an increase in the exponent γ (the network growth is becoming more super-linear). In the cases of Fig. 2.8(a), (b) and (e), the form of non-linear attachment (i.e., sub-linear, $\gamma < 1$, or super-linear, $\gamma > 1$) does not change. However, in the case of the Wikipedia hyperlink network in Fig. 2.8(f), the relative attachment kernel early in the network's evolution appeared super-linear, whereas by its endpoint, the attachment kernel was measured to be sub-linear. Here we reiterate our previous criticism: if the attachment kernel of a network is meant to predict its future evolution, how do we reconcile that measurements across some time windows result in one prediction, while other time windows result in a different, wholly incompatible prediction.

The story in Fig. 2.8(c) and (d) is more complex. In both cases, the ratios of the relative attachment kernels initially appear to decrease below 1, before increasing at moderate degree and exceeding a ratio of 1 at large degree. This behaviour is qualitatively very similar to the dynamics observed in Fig. 2.4 for the k2 model. In these cases, approximating the change in the attachment kernel as a change in the non-linear preferential attachment kernel γ is not easy since the data for small degree may imply an increase in γ whereas the data for large degrees may imply a decrease in the exponent γ. In such cases, simple assumptions of non-linear preferential attachment are not sufficient to draw reliable conclusions about the future evolution or past origin of a network.

A number of mechanisms may be responsible for the appearance of time dependence in the relative attachment kernel. Generally, the simplest explanations for time dependence in network growth relate to changes in network parameters over time. In most network growth models

(including in the k2 model) these are assumed to be time independent for simplicity. An example of such a parameter includes the number of edges added by each new node added to a network, m.

In the BA model, the limiting degree distribution is given by

$$p_\infty(k) = \frac{2m(m+1)}{k(k+1)(k+2)}, \tag{2.29}$$

where we note that this solution is valid for sufficiently large graphs given any initial network at time t_0. As a result, if we let $m \to m(t)$, it is plausible that we will observe transient behaviour during which the form of the degree distribution may change over time. The same argument holds for measurements of the network attachment kernel. Such time dependence in m is not hypothetical and has been shown to be true by Leskovec, Kleinberg, and Faloutsos [53] for a number of different growing networks including citation networks, patent networks and affiliation networks. Some network models consider growth in the average out-degree of nodes over time, see for instance [83]. However, despite showing complex time dependent scaling in the time evolution of individual nodes (the analytical form for the degree distribution is not solvable), the authors argue that the limiting degree evolution implies power law scaling in the degree distribution for large graphs. This may be true, but as the k2 model illustrates, in some cases the transient phase of a network may be so long such that, for all practical purposes, the limiting degree distribution is not necessarily observed during the lifetime of a real world network. Slow convergence to a limiting degree distribution has been noted previously for node copying models in [78].

Another simple parameter which may affect the time dependence of either the degree distribution or attachment kernel of a network is the exponent for non-linear preferential attachment. Consider letting $\gamma \to \gamma(t)$ for $\Pi(k) \propto k^\gamma$. It is already known that γ effects the limiting degree distribution of non-linear preferential attachment [43], and that for $\gamma > 1$, the degree distribution appears scale-free for a transient period [44]. Hence, any variation in γ is likely to add to the complexity of the time dependence observed in network growth observables.

So far, we have provided only two examples of parameters whose time dependence may alter the transient behaviour of a network's degree distribution or attachment kernel. However, we would argue that any solvable network model where a constant parameter appears in the analytical form for the degree distribution is likely to exhibit time dependent transients if that constant becomes time dependent itself.

In cases like the k2 model, where no individual parameter has been set to be time dependent, the reasons for complex scaling in the observables of network growth are less easily explained and may not be easy to elucidate from data. In the k2 model, the origin lies in the implicit time dependence of the local network structure which results in super-linear scaling in the influence network, associated with gelation to important communities. However, any network growth model where the attachment kernel is determined by an observable which implicitly changes over time as the structure of the network changes is likely to exhibit similar time dependence. For instance, network growth based on attaching to nodes according to their betweenness is likely to exhibit a time dependent network attachment kernel, as investigated by [67].

We note to the reader that it was our original intention to perform the analysis in Fig. 2.8 with statistical rigour. However, this task has proven difficult given that (1) our data breaks many of the assumptions underlying common statistical tests, and (2) it is not yet fully understood how techniques for estimating attachment kernels, like Newman's method, are effected/biased by time dependence – by construction, these techniques assume the attachment kernel is time independent. We highlight the need to tackle these problems with better statistical network analysis in future work.

2.5 Discussion

The study of complex networks has come to dominate complexity science in the 21st century, and is likely to become more prominent in a hyper-connected world. Not only have complex networks become influential in physics and mathematics, but their trans-disciplinary appeal has led to their use across almost all areas of science and academia, from archaeology [84] to

neuroscience [85], economics [86] to epidemics [87], and many more.

A key feature of network science is the study of how networks emerge and evolve over time, and numerous models and techniques have been developed to explore this problem [13, 38, 43, 45, 46, 51, 61, 73, 76, 83, 88–92]. In almost all cases, these models and techniques have their limitations and are only applicable to the real world under a number of key constraints. Despite this, the spread of network science has been so extensive that many of these approaches are being used without a robust understanding of their underlying assumptions. In this chapter we have discussed two such assumptions: (1) that the rules underlying network growth do not depend on time, and (2) that the degree of nodes in a network is the key observable determining network evolution.

The number of research papers discussing network growth and attempting to infer their underlying mechanisms is vast, often guided by simple network models to inform their analysis. However, the models most frequently discussed in popular network science textbooks, for instance [38, 73], almost always assume that underlying growth rules are fixed in time, and (as is the case for the BA model, but not the k2 model) that the growth rule itself is stationary [3]. As a result, it is not particularly surprising that most papers inferring network growth mechanisms also make these assumptions. In many contexts such assumptions are sensible and essential, allowing for analytically tractable calculations which may otherwise be impossible. However, in some real world scenarios such approaches may not be suitable. A selection of papers which do consider the implication of these assumptions include [65, 93] in the context of preferential attachment models, [68, 94] in organisational networks, [67] in social networks, and others [62, 63, 66].

In this chapter, we have tried to highlight how very simple network growth rules can break both the time-independence of the network degree distribution and the time-independence of the node-node attachment probability. We have done this by introducing the k2 model, a simple variant of the Barabási-Albert model where the attractiveness of a node is correlated

[3] Note that the k2 model assumes that the time derivative of the growth rule is itself stationary. This may not be the case for real networks but is an acceptable first-order approximation to highlight the risks of assuming that the growth rule does not change at all over time.

to the attractiveness of a node's neighbours. Even though such a network growth rule does not contain an explicit time dependence, the formation of clusters means that a node's attractiveness is implicitly time dependent through its dependence on its local environment. This mechanism is relevant for real-world networks involving mutual benefit where a node gains an advantage from being connected to an influential neighbour, such as in collaboration networks [60], or citation networks [61], or indirectly in systems with neighbour-neighbour interactions and copying processes [62, 63, 70, 75–79].

The k2 model shows that for small networks, the degree distribution appears approximately power law, and the attachment kernel is approximately linear, both of which are consistent with preferential attachment. However, after a lengthy transient period, both the degree distribution and attachment kernel show significant deviations from the simple scaling predicted for preferential attachment. These deviations grow over time showing strong time dependence. We support these findings with an approximate analytical treatment showing that assumptions of simple scaling forms in the evolution of individual nodes in the k2 model results in inconsistencies in the mathematics, suggesting that numerous scaling regimes are interacting and changing over time.

The k2 model is an idealised network growth model – it does not reflect real-world networks, even if the underlying mechanism has explanatory value. However, changes in the degree distribution and the attachment kernel can also be seen in real data of varying origins. In six networks for which dynamic network data is available (three social networks, one hyperlink network, one collaboration network, and one citation network), we have found that these networks are approximately time independent on short timescales, but show significant time dependency, and diversity in that dependency, over longer timescales. In some cases, this time dependency may have simple origins such as node ageing, changes in the average out-degree over time, or changes in the exponent for preferential attachment. However, in some cases, time dependency may arise implicitly resulting in complex scaling.

How to distinguish between these cases using real data is an open question; current methods available for estimating attachment kernels still assume stationary kernels, and it is not clear

how to adapt methods to avoid this assumption [51]. Future work should explore this question, for instance by applying attachment kernel estimation methods to a wider selection of non-stationary growth rules to understand how growth rule time dependence affects the estimated kernel.

2.6 Conclusions

The k2 model is constructed principally for illustrative purposes, highlighting the importance of time dependence in the rules of network growth. Probing a number of real networks, it is clear that, at least on longer time scales, the rules of network growth are certainly time dependent. Therefore, as long as network science techniques are being applied to the real world by experts and non-experts alike, it is essential that we understand the limitations of simple models and consider their underlying assumptions. In many cases the origin of the time dependence may have a simple explanation, which, if accounted for in prediction models, may avoid excessive errors in forecasting the evolution of networks and the dynamics taking place on those networks. However, knowing the impact of these assumptions is only possible if simple steps are taken to check their validity.

In the context of the wider debate on "scale-free" networks, it is worth considering the following. If the generative mechanisms underlying network growth are not constant in time, is it plausible that the functional form for network degree distributions will be constant in time? In many cases the change over time may be very small. However, if we apply a strict definition of "scale-freeness," small changes in the degree distribution may be sufficient to induce changes in the most-likely functional form for the degree distribution as predicted using the current state of the art measures [19]. This may explain why only 4% of real world networks have been identified as scale-free [19].

2.7 Chapter summary

- Most discussions of real-world network structure do not discuss the time dependence of network degree distributions.
- Similarly, most growing network models assume time independent growth rules.
- We introduce the k2 model, a network growth model based on linear preferential attachment proportional to a node's second degree. This rule has an implicit time dependence, resulting in time dependent degree distributions.
- To analyse the k2 model we decompose the network into an observed and influence layer (the influence layer is referred to as the hidden layer in the next chapter). Bidirectional coupling between these two layers is, in part, responsible for the emergent time dependence in the k2 model.
- Drawing on the results for the k2 model, we identify time dependent network growth rules in a range of real-world networks.
- We suggest that future work must make a greater effort to discuss time dependence in network growth, particularly when discussing the past or future states of a system.

2.8 Mathematical derivations

2.8.1 Derivation of Eq. (2.1).

When $m = 1$, it is convenient to write the first and second degree of node j, using the adjacency matrix $a_{i,j}(t)$, which encodes direct connections in our network. Here, $a_{i,j}(t) = 1$ if an edge exists between nodes i and j in the observed network, $a_{i,j}(t) = 0$ otherwise. In the adjacency matrix formulation, the first degree of node j can be written as

$$k_j^{(1)}(t) = \sum_{i=1}^{N(t)} a_{i,j}(t), \tag{2.30}$$

and the second degree of node j can be written as

$$k_j^{(2)}(t) = \sum_{i=1}^{N(t)} a_{i,j}(t) \cdot k_i^{(1)}(t). \tag{2.31}$$

Subbing Eq. (2.30) into Eq. (2.31) and summing over all the nodes in the network we can write

$$\sum_{j=1}^{N(t)} k_j^{(2)}(t) = \sum_{j=1}^{N(t)} \sum_{i=1}^{N(t)} a_{i,j}(t) \cdot k_i^{(1)}(t). \tag{2.32}$$

In the case of $m = 1$ where the tree structure of the network allows us to ignore cross-terms, we can swap the order of the summations such that

$$\sum_{j=1}^{N(t)} k_j^{(2)}(t) = \sum_{i=1}^{N(t)} k_i^{(1)}(t) \cdot \sum_{j=1}^{N(t)} a_{i,j}(t) = \sum_{i=1}^{N(t)} (k_i^{(1)}(t))^2, \tag{2.33}$$

as required. Note, however, that this argument is strictly only valid if $m = 1$. If $m > 1$, Eq. (2.31) does not give the number of nodes distance 2 from node j, but rather gives the number of paths of length two originating from node j. In a tree, these definitions are equivalent. However, in a network with loops the number of nodes up to two steps away from node j is strictly no larger than the number of paths of length 2. Hence, we avoid the adjacency matrix notation in other parts of this work where it is not explicitly necessary for derivation purposes.

For $m > 1$, we can test the impact of loops on the validity of Eq. (2.33). Figure 2.9 plots the ratio of the two sums, for $m = 1, 3$, defined as

$$\frac{S_2(t)}{S_1(t)} = \frac{\sum_{j=1}^{N} k_j^{(2)}(t)}{\sum_{l=1}^{N} (k_l^{(1)}(t))^2}, \tag{2.34}$$

against time. The figure has been averaged over 100 simulations of the k2 model. Figure 2.9 indicates that for $m = 1$, $S_2(t)/S_1(t) = 1$ for all t, as expected. For $m > 1$, there is a noticeable difference between $S_2(t)$ and $S_1(t)$ at very small times in the network's evolution. This is to be expected since when the network is small, the probability of acquiring non-unique second degree neighbours is small, but not negligible. As the network evolves, the ratio $S_2(t)/S_1(t)$

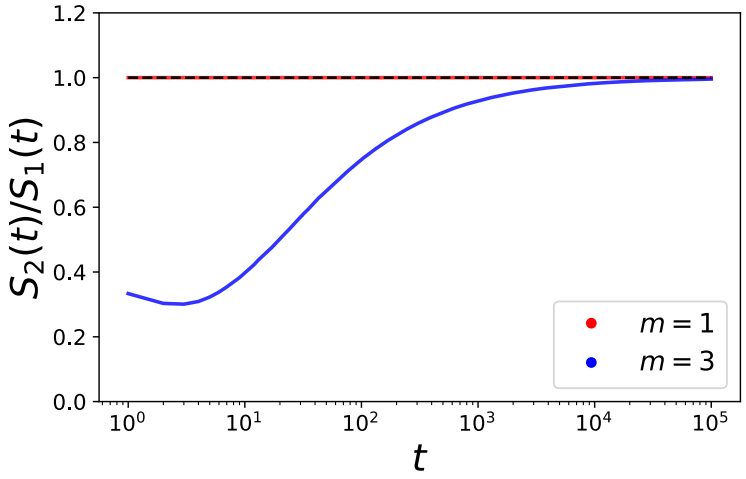

Figure 2.9: The ratio of the sum over the second degree of each node in the network to the sum over the first degree squared for each node indicated in the network, see Eq. (2.34). The ratio equals one for $m = 1$, and converges to one for $m > 1$.

quickly converges to 1, with $S_2(t)/S_1(t) > 0.9$ by $t = 10^3$. This indicates that Eq. (2.34) is a good approximation even for $m > 1$.

2.8.2 Derivation of Eq. (2.10).

Let us consider the node i, the neighbours i_α connected to node i, and the node $i_{\alpha(\beta)}$ connected to node i_α. Here, $\alpha = 1, \cdots, k_i^{(1)}(t)$, and $\beta = 1, \cdots, k_\alpha^{(1)}(t) - 1$. Then, the k2 model with $m = 1$ is applied to nodes i_α as follows:

$$\frac{d}{dt} k_{i_\alpha}^{(1)}(t) = \frac{k_{i_\alpha}^{(2)}(t)}{\sum_{j=1}^{N} k_j^{(2)}(t)} \quad \text{for} \quad \forall \alpha. \tag{2.35}$$

Next, by using Eq. (2.7), the numerator of the right side of Eq. (2.35) can be rewritten as

$$k_{i_\alpha}^{(2)}(t) = k_i^{(1)}(t) + \sum_{\beta=1}^{k_\alpha^{(1)}(t)-1} k_{i_{\alpha(\beta)}}^{(1)}(t). \tag{2.36}$$

Therefore, we obtain the dynamical equations of node i_α as follows:

$$\frac{d}{dt} k_{i_\alpha}^{(1)}(t) = \frac{k_i^{(1)}(t) + \sum_{\beta=1}^{k_\alpha^{(1)}(t)-1} k_{i_{\alpha(\beta)}}^{(1)}(t)}{\sum_{j=1}^{N} \left(k_j^{(1)}(t) \right)^2}. \tag{2.37}$$

Summing $k_{i_\alpha}^{(1)}(t)$ from $\alpha = 1$ to $\alpha = k_i^{(1)}(t)$ for each side, we obtain

$$\frac{d}{dt}\left(\sum_{\alpha=1}^{k_i^{(1)}(t)} k_{i_\alpha}^{(1)}(t)\right) = \frac{\sum_{\alpha=1}^{k_i^{(1)}(t)}\left(k_i^{(1)}(t) + \sum_{\beta=1}^{k_\alpha^{(1)}(t)-1} k_{i_\alpha(\beta)}^{(1)}(t)\right)}{\sum_{j=1}^{N}\left(k_j^{(1)}(t)\right)^2}. \tag{2.38}$$

Recall that, $k_i^{(2)}(t) = \sum_{\alpha=1}^{k_i^{(1)}} k_{i_\alpha}^{(1)}(t)$. Thus, Eq. (2.38) can be modified as follows:

$$\begin{aligned}
\frac{d}{dt}k_i^{(2)}(t) &= \frac{\sum_{\alpha=1}^{k_i^{(1)}(t)} k_i^{(1)}(t) + \sum_{\alpha=1}^{k_i^{(1)}(t)} \sum_{\beta=1}^{k_\alpha^{(1)}(t)-1} k_{i_\alpha(\beta)}^{(1)}(t)}{\sum_{j=1}^{N}\left(k_j^{(1)}(t)\right)^2} \\
&= \frac{k_i^{(1)}(t)k_i^{(1)}(t) + \sum_{\alpha=1}^{k_i^{(1)}(t)} \sum_{\beta=1}^{k_\alpha^{(1)}(t)-1} k_{i_\alpha(\beta)}^{(1)}(t)}{\sum_{j=1}^{N}\left(k_j^{(1)}(t)\right)^2} \\
&= \frac{\left(k_i^{(1)}(t)\right)^2 + \sum_{\alpha=1}^{k_i^{(1)}(t)} \sum_{\beta=1}^{k_\alpha^{(1)}(t)-1} k_{i_\alpha(\beta)}^{(1)}(t)}{\sum_{j=1}^{N}\left(k_j^{(1)}(t)\right)^2}.
\end{aligned} \tag{2.39}$$

This is equivalent to Eq. (2.14).

Chapter 3

Heterogeneous node copying from hidden network structure

3.1 Publication

Chapter 3 is based on the following publication:

M. Falkenberg. "Heterogeneous node copying from hidden network structure". In: *Communications Physics* 4.1 (2021), p. 200. DOI: 10.1038/s42005-021-00694-1

Contributions: The work in this chapter is wholly mine, although I must credit a number of invaluable discussions with Chester Tan and Tim Evans without whom this work would not have been initiated. In particular, I am indebted to Tim Evans for pointing out similarities between my work in chapter 2 and the copying model in [62] which inspired this work.

3.2 Introduction

Node copying is an important network growth mechanism [62, 95–100]. In social networks, copying is synonymous with triadic closure, playing an important role in the emergence of high

clustering [101, 102]. In biology, node copying encapsulates duplication and deletion, a key mechanism in the formation of protein-interaction networks [90, 103–106].

Despite this range of applications, most node copying models assume uniform, or homogeneous copying, i.e., that the probability of copying any given neighbour of a node is equal. The exact formulation varies widely, but examples include "links are attached to neighbours of [node] j with probability p" [98], or "one node [is duplicated]... edges emanating from the newly generated [node] are removed with probability δ" [104]. Many other models use similar uniform copying rules [62, 76, 90, 96, 99–101, 103, 107–117].

Homogeneous copying is a sensible base assumption, often aiding a model's analytical tractability. However, especially in a social context, there is good reasons to believe that node copying may be heterogeneous. As an example, consider the social brain hypothesis, a theory which suggests that the average human has around 150 friends (Dunbar's number), encapsulating progressively smaller sub-groups of increasing social importance [118, 119]. In contrast, large social networks often have an average degree far exceeding Dunbar's number [120], implying that most of these observed friends are only distant acquaintances. In this context, if individual A introduces individual B to one their friends, C, (i.e., B is copying A's friend C), we may reasonably expect that C is more likely to be chosen from A's inner social circle, than A's wider social circle.

This is directly related to the principle of strong triadic closure: "If a node A has edges to nodes B and C, then the B-C edge is especially likely to form if A's edges to B and C are both strong ties" [121]. In weighted networks where tie strength can be equated to edge weight, empirical evidence for the strong triadic closure principle can be inferred by measuring the neighbourhood overlap between two nodes as a function of tie strength [121]; for example using mobile communication networks [122], or using face to face proximity networks [123].

Unfortunately, for many networks tie strength data is unavailable or unknown; such data may be available to some parties, e.g., the social media company, but is often not publicly available. In these cases, evidence for asymmetric triadic closure may be inferred through proxy means. For instance, in academic collaboration networks it has been shown that the ratio of triadic closure

varies strongly with the number of shared collaborators between nodes [124], without requiring any knowledge as to how many papers were co-authored by any two individuals. Although the average triadic closure ratio is small (typically $< 10\%$), the ratio rapidly increases with the number of shared collaborators. However, these aggregate measures are highly coarse grained and likely only approximate real closure dynamics.

This motivates the study of simple heterogeneous copying models [97, 98, 102, 125, 126]. Typically these models fall into a small number of distinct categories. In the first, heterogeneity is introduced as a node intrinsic property (e.g., node fitness) in the absence of structural considerations [98]. In the second, heterogeneity is introduced via group homophily where the probability of triadic closure between nodes A and B is dependent on whether nodes A and B are in the same group or different groups (e.g., researchers from the same academic discipline, as opposed to different disciplines) [102, 126]. However, intra-group copying is typically modelled uniformly. Finally, some models consider heterogeneous copying driven by the network structure around nodes A and B, without introducing node homophily [97].

Bhat, Krapivsky, and Redner [97] define a threshold model where node A introduces node B to one of their friends C. An edge then forms between B and C if the fraction of neighbours common to B and C exceeds some threshold F. The model demonstrates a transition from a state where networks are almost complete for small F, to a state where networks are sparse but highly clustered as F increases past a critical threshold. However, the model is limited in its tractability and has peculiarities such as the observation that fringe communities are almost always complete.

In the current work, our aim is to extend these ideas and introduce a more general framework for heterogeneous node copying based on the concept of hidden strong ties. To do so, we introduce the hidden network model, a framework based on multilayer networks [17] where layers have identical node structure but different edge structure. The framework lets us build models where local heterogeneity in the rules of network growth is a property of the hidden network structure and not arbitrarily encoded using node intrinsic properties or group homophily. The concept is closely related to other multilayer paradigms including the use of replica nodes to model

heterogeneity [127], interdependent networks [128], and multilayer copying [113].

In the remainder of this chapter, we define and analytically study the case of extreme heterogeneous copying, the correlated copying model (CCM). The CCM is an adaptation of the uniform copying model introduced by Lambiotte *et al.* [62], see Fig. 3.1(a). In the UCM, a single node, α, is added to the network at time t_α, and connects to one target node, β, which is chosen uniformly at random. The formation of an edge between the new node and the target node puts the UCM in the class of corded copying models; Steinbock, Biham, and Katzav [108] refer to the UCM as the corded node duplication model. We label each neighbour of β with the index γ_j where $j \in \{1, \cdots, k^\beta\}$, and k^β is the degree of node β. For each neighbour γ_j, the copied edge (α, γ_j) is added to the network independently with probability p.

Following the convention of previous copying models, the nodes α and β are sometimes referred to as the daughter and mother nodes respectively. The network is initialised at $t = 1$ with a single node. If $p = 0$, no edges are copied resulting in a random recursive tree. If $p = 1$, the UCM generates a complete graph.

In addition to the extreme copying case (the CCM), we numerically investigate a generalised form of the correlated copying model (GCCM) which interpolates between the UCM and CCM. The GCCM generates a diverse spectrum of network structures spanning both ergodic sparse and non-ergodic dense networks, with degree distributions ranging from exponential decay, through stretched-exponentials and power laws, to extremely fat tailed distributions with anomalous fluctuations. These networks exhibit a broad clustering spectrum from sparse networks with significantly higher clustering than their uniform equivalents, to the unusual case where networks are almost complete, but with near zero clustering. We comment on a selection of real collaboration networks, which, in line with the CCM, exhibit higher clustering than can be explained by uniform copying. This suggests that heterogeneous copying may be an important explanatory mechanism for social network formation.

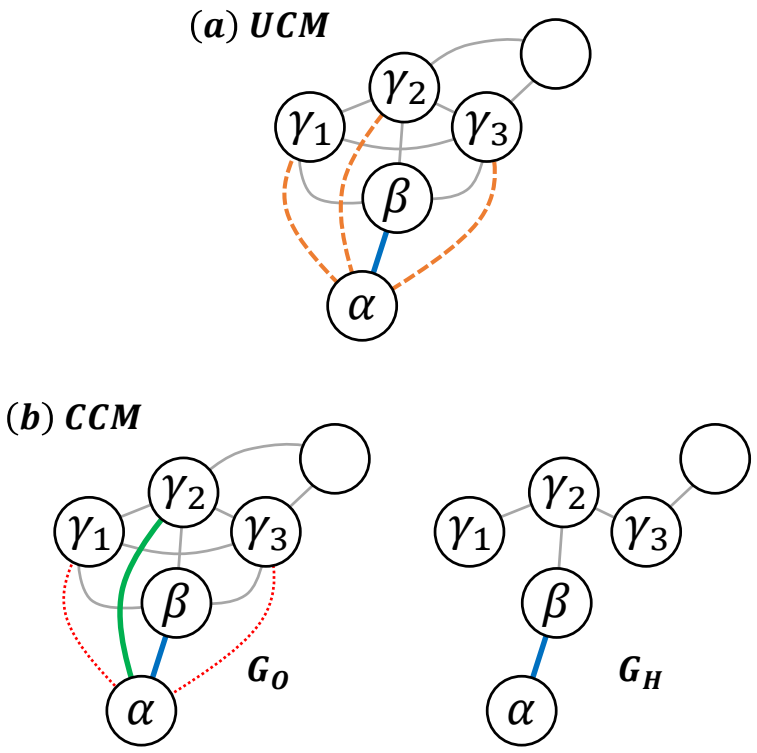

Figure 3.1: Two models of network formation via node copying. (a) The uniform copying model (UCM), and (b) the correlated copying model (CCM). The UCM consists of a single layer. The CCM has an observed layer, in which copying takes place, G_O, and a hidden layer, G_H. For both models, a new node α (the daughter) is added to the existing network (nodes connected by grey edges) and forms a random link (blue) to a target node, β (the mother). (a) In the UCM, there is a uniform probability, p, of forming an edge to each of β's neighbours ($\gamma_1, \gamma_2, \gamma_3$; orange dashed edges). (b) In the CCM, copied edges are added to the observed network, G_O, deterministically. If an edge exists in the hidden network, G_H, between node β and node γ_j (e.g., the $\{\beta, \gamma_2\}$ edge), then node α copies that edge in G_O (e.g., forming the $\{\alpha, \gamma_2\}$ edge; solid green). If an edge does not exist in G_H (e.g., the $\{\beta, \gamma_1\}$ and $\{\beta, \gamma_3\}$ edges), the corresponding edges are not copied to G_O (red dotted lines). Copied edges are never added to G_H, such that G_H is always a tree; this is necessary to ensure the analytical tractability of the CCM

.

3.3 Correlated copying model

The correlated copying model is constructed using a hidden network model, defined as the pair of single layer graphs $G = (G_O, G_H)$, comprising an observed network $G_O = (V, E_O)$ and a hidden network $G_H = (V, E_H)$, where V is the set of nodes for both networks and E_O and E_H are the set of edges for each network. The set V represents the same entities in G_O and G_H, with differences lying exclusively in the edge structure between nodes.

In the CCM, see Fig. 3.1(b), the observed and hidden networks are initialised with a single

node at $t = 1$. At $t = t_\alpha$, node α is added to both networks and a single target node, β, is chosen uniformly at random. We label the k_O^β neighbours of β in G_O with the index γ_j. Then, in the observed network only, the copied edge (α, γ_j) is formed with $p_{\text{hid}} = 1$ if the edge $(\beta, \gamma_j) \in E_H$, $p_{\text{obs}} = 0$ otherwise. The general case with intermediate copying probabilities is discussed in section "General correlated copying model". No copied edges are added to the hidden network, G_H. The direct edge (α, β) is added to both G_O and G_H. The CCM therefore also falls into the class of corded node duplication models. Using the convention of referring to β as the mother node and α as the daughter node, we note that the hidden network consists exclusively of first-order relations (mother-daughter), whereas edges found only in the observed network correspond to second-order relations (sister-sister, or grandmother-granddaughter).

G_H evolves as a random recursive tree. Unlike the UCM, all copying in G_O is deterministic, with the only probabilistic element emerging in the choice of the target node β. For comparative purposes, we define the effective copying probability in the CCM as $p_{\text{eff}} = \langle k_H^\beta / k_O^\beta \rangle$, i.e., the fraction of the observed neighbours of node β which are copied by node α, averaged over time.

Framed in a social context, we might think of G_O as an observed social network where individuals have many friends, but the quality of those friendships is unknown, with most ties being weak. In contrast, underlying every social network is a hidden structure representing the inner social circle of individuals, where a node is only connected to their closest friends [119]. Copying in the CCM is biased to this inner circle.

3.3.1 Basic topological properties

The total number of edges in G_H is given by $E_H(t) = t - 1$, with the average degree given by $\langle k_H \rangle \approx 2$. Using the degree distribution of G_H, see below, $\langle k_H^2 \rangle = 6$. In the observed network, each time step a single edge is added by direct attachment, and one copied edge is added for each neighbour of the target node in the hidden network, G_H, which is written as the hidden degree of node β, k_H^β. The average change in the number of edges is therefore $\langle \Delta E_O(t) \rangle = 1 + \langle k_H^\beta \rangle = 1 + \langle k_H \rangle = 3$, such that $\langle E_O(t) \rangle \sim 3t$ and $\langle k_O \rangle = 6$.

As an alternative, note that the hidden degree of node α can be written as the sum over row α of the adjacency matrix, $a_{\alpha,\beta}$,

$$(k_H)_\alpha = \sum_{\beta=1}^{t} a_{\alpha,\beta}. \tag{3.1}$$

Similarly, the observed degree of node α can be written as

$$(k_O)_\alpha = \sum_{\beta=1}^{t} a_{\alpha,\beta}(k_H)_\beta \tag{3.2}$$

into which we can substitute Eq. (3.1) and sum over all α to give

$$\sum_{\alpha=1}^{t}(k_O)_\alpha = \sum_{\alpha=1}^{t}\sum_{\beta=1}^{t} a_{\alpha,\beta}(k_H)_\beta = \sum_{\beta=1}^{t}(k_H)_\beta \cdot \sum_{\alpha=1}^{t} a_{\alpha,\beta} = \sum_{\beta=1}^{t}(k_H)_\beta^2, \tag{3.3}$$

where after the second equality we have changed the order of the sums which is possible due to the absence of cross-terms (since the hidden network is a tree). Averaging both sides of Eq. (3.3) over all t nodes we find,

$$\langle k_O \rangle = \frac{1}{t}\sum_{\alpha=1}^{t}(k_O)_\alpha = \frac{1}{t}\sum_{\beta=1}^{t}(k_H)_\beta^2 = \langle k_H^2 \rangle. \tag{3.4}$$

This calculation is identical to Eq. (2.5) in chapter 2. In section 3.9, a similar argument is used to derive $\langle k_O^2 \rangle \approx 62$.

We may naively expect that the effective copying probability is $p_{\text{eff}} = \langle k_H \rangle / \langle k_O \rangle = 1/3$. However, for the CCM, $p_{\text{eff}} = \langle k_H^\beta / k_O^\beta \rangle \neq \langle k_H \rangle / \langle k_O \rangle$. We have not found a route to calculating this exactly, but simulations suggest $p_{\text{eff}} \approx 0.374$.

3.3.2 Degree distribution

Here, our aim is to derive the degree distribution for the observed network analytically. The hidden network evolves as a random recursive tree which has a limiting degree distribution given by

$$p_H(k_H) = 2^{-k_H}, \text{ for } k_H > 1. \tag{3.5}$$

In section 3.9, we show that the degree distribution for the observed network can be written as the recurrence

$$p_O(k_O) = \frac{\pi_O(k_O - 1) \cdot p_O(k_O - 1) + 2^{1-k_O}}{1 + \pi_O(k_O)}, \text{ for } k_O \geq 2, \tag{3.6}$$

where the final term is the probability that at time t the newly added node has initial degree k_O and

$$\pi_O(k_O) = 1 + \langle k_H \mid k_O \rangle, \tag{3.7}$$

with $\langle k_H \mid k_O \rangle$ as the average degree of nodes in the hidden network with observed degree k_O. Here, the 1 corresponds to edges that are gained from direct attachment, whereas $\langle k_H \mid k_O \rangle$ corresponds to edges gained from copying. Although we have not found an exact expression for $\langle k_H \mid k_O \rangle$, we can make progress by considering the evolution of individual nodes.

Consider node α added to the network at time t_α. The initial conditions for node α are

$$\begin{aligned} (k_H(t_\alpha))_\alpha &= 1, \\ \langle k_O(t_\alpha) \rangle_\alpha &= 1 + \langle k_H(t_\alpha - 1) \rangle_\beta, \end{aligned} \tag{3.8}$$

where the final term is the average hidden degree of the target node β. In G_H, node α gains edges from direct attachment only. Hence, at $t > t_\alpha$,

$$\langle k_H(t) \rangle_\alpha = 1 + \sum_{j=t_\alpha}^{t-1} \frac{1}{j} = 1 + H_{t-1} - H_{t_\alpha - 1}, \tag{3.9}$$

where H_n is the n^{th} harmonic number. In G_O, either node α is targeted via direct attachment, or a copied edge is formed from the new node to node α via any of the $(k_H(t))_\alpha$ neighbours of node α. Hence,

$$\begin{aligned} \langle k_O(t) \rangle_\alpha &= \langle k_O(t_\alpha) \rangle_\alpha + \sum_{j=t_\alpha}^{t-1} \frac{1 + \langle k_H(j) \rangle_\alpha}{j} \\ &= \langle k_O(t_\alpha) \rangle_\alpha + \sum_{j=t_\alpha}^{t-1} \frac{2 + H_j - H_{t_\alpha - 1} - 1/j}{j}, \end{aligned} \tag{3.10}$$

where we have subbed in Eq. (3.9) and $H_{j-1} = H_j - 1/j$. Evaluating this sum, see section 3.9,

we find

$$\langle k_O(t)\rangle_\alpha = \langle k_O(t_\alpha)\rangle_\alpha + \frac{1}{2}\left[(4 + H_{t-1} - H_{t_\alpha-1})(H_{t-1} - H_{t_\alpha-1}) - H^{(2)}_{t-1} + H^{(2)}_{t_\alpha-1}\right], \tag{3.11}$$

where $H^{(m)}_n$ is the n^{th} generalised Harmonic number of order m. For $t \to \infty$, $H^{(2)}_t \to \pi^2/6$. Hence, for large t we can drop the final two terms and substitute in Eq. (3.9) to give

$$\langle k_O(t)\rangle_\alpha \approx \langle k_O(t_\alpha)\rangle_\alpha + \frac{1}{2}\left(\langle k_H(t)\rangle_\alpha + 3\right)\left(\langle k_H(t)\rangle_\alpha - 1\right). \tag{3.12}$$

Noting, that Eq. (3.12) is a monotonically increasing function of k_H for $k_H > 1$, we assume that we can drop the index α and the time dependence giving the average observed degree of nodes with specific hidden degree as

$$\langle k_O \mid k_H\rangle \approx \langle \tilde{k}_O \mid k_O\rangle + \frac{1}{2}\left(k_H + 3\right)\left(k_H - 1\right), \tag{3.13}$$

where $\langle \tilde{k}_O \mid k_O\rangle$ denotes the average initial observed degree of nodes with current degree k_O. Finally, we make the approximation that $\langle k_H \mid k_O\rangle \approx \langle k_O \mid k_H\rangle^{-1}$ where the exponent denotes the inverse function. This gives

$$\pi_O(k_O) = 1 + \langle k_H \mid k_O\rangle \approx \sqrt{2(k_O + 2 - \langle \tilde{k}_O \mid k_O\rangle)}. \tag{3.14}$$

To proceed, let us solve the degree distribution at $k_O = 2$. Although the average initial condition $\langle \tilde{k}_O\rangle = 1 + \langle k_H\rangle = 3$, in this case $\langle \tilde{k}_O \mid 2\rangle = 2$. Therefore

$$p_O(2) = -\pi_O(2)\cdot p_O(2) + 2^{-1} = -p_O(2)\cdot\sqrt{2(2)} + 2^{-1}, \tag{3.15}$$

giving $p_O(2) = 1/6$. Since $\langle \tilde{k}_O \mid k_O\rangle$ has an almost negligible effect on $\pi_O(k_O)$ for $k_O > 2$, for simplicity we set $\langle \tilde{k}_O \mid k_O\rangle = 2$. We can now rewrite Eq. (3.6) as

$$p_O(k_O) = \frac{p_O(k_O - 1)\sqrt{2(k_O - 1)} + 2^{1-k_O}}{1 + \sqrt{2k_O}}, \text{ for } k_O > 2. \tag{3.16}$$

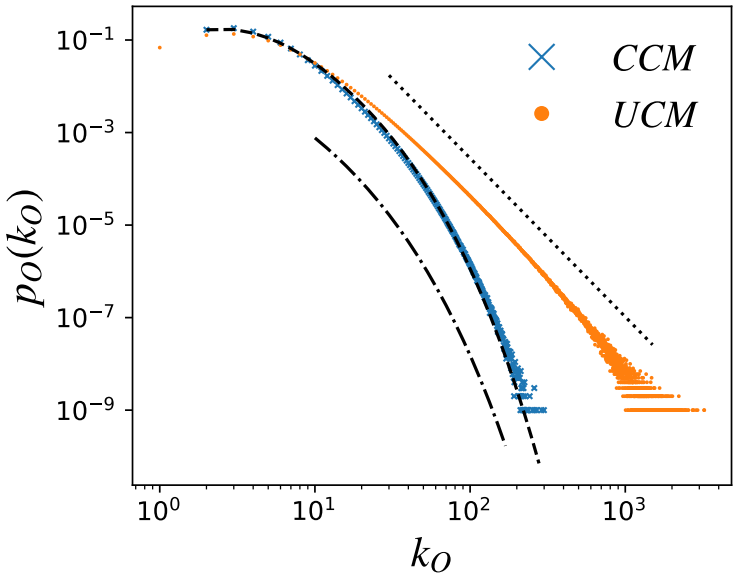

Figure 3.2: The degree distributions for the correlated copying model (CCM) and uniform copying model (UCM). Degree probability, $p_O(k_O)$, plotted as a function of the observed degree, k_O. UCM initialised with copying probability $p = 0.374$ (equal to the CCM's effective copying probability). Networks grown to $t = 10^7$, averaged over 100 networks. Error bars omitted for clarity. Dashed line: Analytical expression for CCM in Eq. (3.16). Dot-dashed: stretched exponential approximation. Dotted: power law scaling.

Although computing this recurrence shows good agreement with simulations, see Fig. 3.2, we have not found a closed form solution to Eq. (3.16).

As an approximation, we return to Eq. (3.11) and note that $H_{t-1} - H_{t_\alpha - 1} \approx \ln(t/t_\alpha)$. Substituting this into Eq. (3.11) and dropping small terms

$$\langle k_O(t > t_\alpha)\rangle_\alpha \approx 2\ln(t/t_\alpha) + \frac{\ln^2(t/t_\alpha)}{2}, \tag{3.17}$$

which inverted gives

$$\ln(t/t_\alpha) \approx -2 + \sqrt{2(k_O+2)} \approx \sqrt{2k_O}, \text{ for } k \gg 2. \tag{3.18}$$

We have dropped the expectation value and define t_α as the time a node was created such that its degree at time t is approximately k_O. Exponentiating each side and taking the reciprocal,

$$\frac{t_\alpha}{t} \approx e^{-\sqrt{2k_O}}. \tag{3.19}$$

Finally, by substituting this approximation into the cumulative degree distribution we find

$$\tilde{p}_O(k_O) = \sum_{k'_O=2}^{k_O} p_O(k'_O) \approx 1 - \frac{t_\alpha}{t} \approx 1 - e^{-\sqrt{2k_O}}, \tag{3.20}$$

which corresponds to a Weibull (stretched exponential) distribution, suppressing the known power law scaling observed in the UCM, see Fig. 3.2 and Ref. [62].

The approximation for the cumulative degree distribution stems from the observation that, on average, nodes with $k'_O > k_O$ were added to the network at $t' < t_\alpha$, whereas nodes with $k'_O < k_O$ were added to the network at $t' > t_\alpha$. Both Eq. (3.18) and Eq. (3.20) are close to the scaling expected from sub-linear preferential attachment [43] with an exponent 1/2.

3.3.3 Clique distribution

In a simple undirected graph, a clique of size n is a subgraph of n nodes which is complete. A clique of size $n = 2$ is an edge, whereas $n = 3$ is a triangle. Here we calculate the exact scaling for the number of n cliques, $Q_n(t)$, in G_O.

Let us first consider the case of triangles. At $t = t_\alpha$, there are two mechanisms by which a new triangle forms:

1. **Direct triangles.** The new node, α, forms a direct edge to the target node, β, and forms copied edges to each of the k_H^β neighbours of node β, labelled with the index γ_j. The combination of the direct edge (α, β), the copied edge (α, γ_j), and the existing edge (β, γ_j) creates one triangle, $(\alpha, \beta, \gamma_j)$, for each of the k_H^β neighbours.

2. **Induced triangles.** If node α forms copied edges to both node γ_j, and to node $\gamma_{j'}$, $j \neq j'$, the triangle $(\alpha, \gamma_j, \gamma_{j'})$ is formed if $(\gamma_j, \gamma_{j'}) \in E_O$.

Combining these mechanisms, the change in the number of triangles can be written as

$$\Delta Q_3(t_\alpha) = \Delta Q_3^D(t_\alpha) + \Delta Q_3^I(t_\alpha), \tag{3.21}$$

where the first and second terms on the right correspond to direct and induced triangles respectively. One new direct triangle is formed for each of the k_H^β neighbours of node β, $\Delta Q_3^D = k_H^\beta$. For induced triangles, the copied edge (α, γ_j) is only formed if $(\beta, \gamma_j) \in E_H$. Additionally, all pairs of nodes which are next-nearest neighbours in G_H must be nearest neighbours in G_O. Hence, the edge $(\gamma_j, \gamma_{j'})$ must exist in the observed network if both γ_j and $\gamma_{j'}$ are copied. As a result, one induced triangle is formed for each pair of copied edges (α, γ_j) and $(\alpha, \gamma_{j'})$ such that

$$\Delta Q_3^I = \binom{k_H^\beta}{2} = \frac{(k_H^\beta)^2 - k_H^\beta}{2}. \tag{3.22}$$

A visual example of the combinatorics for $k_H^\beta = 3$ is shown in Fig. 3.3.

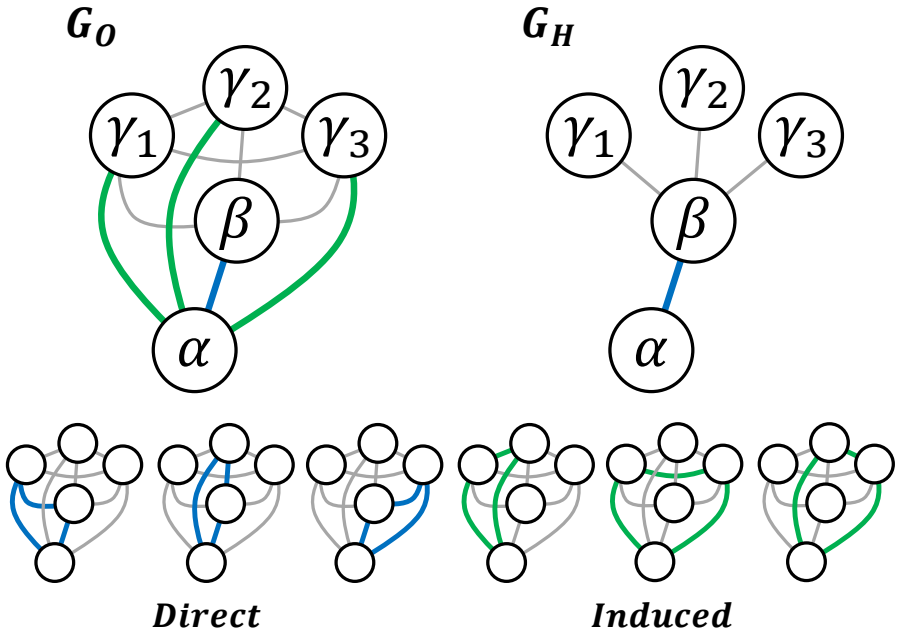

Figure 3.3: A schematic illustrating the number of triangles formed in a single time step of the correlated copying model. G_O: the observed network, G_H: the hidden network. The new node, α, forms a direct edge (blue) to a node β which has three existing hidden neighbours $(\gamma_1, \gamma_2, \gamma_3)$. The copying process forms three new edges (green) in G_O. The copying process results in three new direct triangles (outlined in blue), involving the edge $\{\alpha, \beta\}$, and three new induced triangles (outlined in green), excluding the edge $\{\alpha, \beta\}$. Triangles are formed in the observed network only; the hidden network remains a random tree. Note that in G_O and G_H, the grey edges represent the existing network, while in the small graphs at the bottom they represent edges that are not part of the outlined triangles.

Extending the triangle argument to general n we can write

$$\Delta Q_n(t_\alpha) = \Delta Q_n^D(t_\alpha) + \Delta Q_n^I(t_\alpha), \tag{3.23}$$

where direct cliques are those which include the edge (α, β). For a clique of size n, the number of direct cliques is given by the number of ways in which $n-2$ nodes can be chosen from k_H^β nodes,

$$\Delta Q_n^D(t_\alpha) = \binom{k_H^\beta}{n-2}, \tag{3.24}$$

whereas the number of induced cliques is given by the number of ways in which $n-1$ nodes can be chosen,

$$\Delta Q_n^I(t_\alpha) = \binom{k_H^\beta}{n-1}. \tag{3.25}$$

As $t \to \infty$, the average change in clique number is

$$\langle \Delta Q_n(t) \rangle = \sum_{k_H=1}^{\infty} p_H(k_H) \left[\binom{k_H}{n-2} + \binom{k_H}{n-1} \right], \tag{3.26}$$

where $p_H(k_H)$ is the probability that the randomly chosen target node $k_H^\beta = k_H$. To avoid ill-defined binomials, we rewrite Eq. (3.26) as

$$\langle \Delta Q_n(t) \rangle = p_H(n-2) + \sum_{k_H=n-1}^{\infty} p_H(k_H) \cdot \binom{k_H+1}{n-1}, \tag{3.27}$$

where we have combined the two terms into a single binomial. After substituting in $p_H(k_H)$ and solving the sum,

$$\langle \Delta Q_n(t) \rangle = 2^{2-n} + \sum_{k_H=n-1}^{\infty} 2^{-k_H} \cdot \binom{k_H+1}{n-1} = 4. \tag{3.28}$$

Consequently, for large t we find the curious result that the number of n cliques scales as

$$Q_n(t) \sim 4t, \text{ for } n > 2, \tag{3.29}$$

independent of the clique size. In practice this result only applies for $t \to \infty$. To see this, note that the largest clique in G_O at time t is always directly related to the largest degree node in G_H,

$$\mathrm{Max}(n, t) = \mathrm{Max}(k_H, t) + 1, \tag{3.30}$$

with the largest hidden degree at time t scaling as approximately

$$\mathrm{Max}(k_H, t) \sim \ln(\mathrm{t}). \tag{3.31}$$

We can invert this and ask how large the network is if we observe that the largest observed clique is n. This gives

$$t_n \sim e^{\mathrm{n}}. \tag{3.32}$$

Hence, the scaling relation in Eq. (3.29), is only valid for cliques of size n when $t \gg t_n$.

In Fig. 3.4(a), we plot the average number of cliques of size n for various n as a function of time t, averaged over 1000 simulations. Panel (b) shows the standard deviation of the number of cliques of size n across the same set of 1000 simulations.

The figure shows that for small n, the average number of cliques quickly converges to our expected theoretical value. However, for larger n, the number of cliques has not converged to our expected result by the end of the simulation. Larger clique sizes also show significantly larger fluctuations in the number of cliques than smaller clique sizes. This is expressed by the standard deviation in panel (b) which demonstrates that for large n, the standard deviation in the number of cliques of size n can be many times larger than the average number of cliques, although the standard deviations fall after the initial rapid increase.

To highlight these fluctuations more explicitly, Fig. 3.5 shows three individual simulations for $n = 3$ and three simulations for $n = 10$.

These results support the scaling predicted by theory, although we note that in practice only small clique sizes converge to the expected scaling in timescales which are feasible to simulate numerically.

The predicted clique scaling for the CCM is particularly unusual if considered in the context of random graph theory. Following the work of Erdős and Rényi [129], it is known that the critical occupation probability for observing a clique of size n in a random graph of size t scales as $p_c(t) \sim t^{-2/(n-1)}$. As shown in section 3.3.1, the number of edges for the observed network

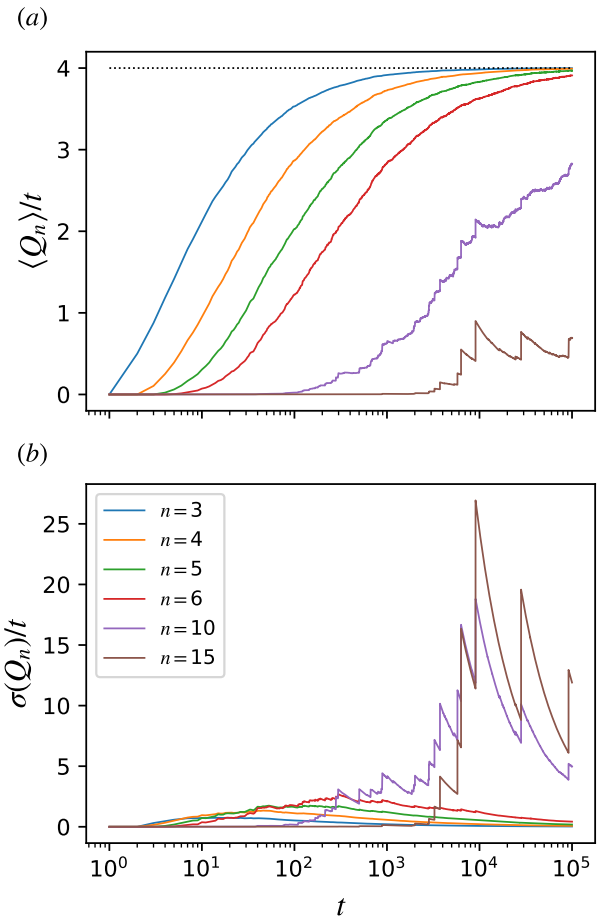

Figure 3.4: Numerical results for the number of cliques of size n rescaled by time t, as a function of t. (a) The mean number of cliques of size n divided by the number of nodes in the network t, averaged over 1000 simulations. Our theoretical result suggests that this result should converge to $\langle Q_n \rangle / t \to 4$ (dotted black line). (b) The rescaled standard deviation for the number of cliques of size n. This demonstrates that there are very large fluctuations in the number of cliques over different simulations.

in the CCM scales as $E_O(t) \sim 3t$, whereas the total number of edges in a complete graph of size t scales as t^2. Hence, the CCM is comparable to a random graph with an edge occupation probability given by $E_O(t)/t^2 \sim t^{-1}$, sitting at the critical edge occupation probability for triangles, $n = 3$; i.e., we do not expect to find any cliques of size $n \geq 3$ as $t \to \infty$. However, in the CCM we observe that we have an equal number of cliques of all sizes (on average) as $t \to \infty$.

This result appears unintuitive at first, but can be understood if we consider a triadic closure on a wedge in both a random graph and the CCM. First, consider a random graph in the limit as $t \to \infty$ and select any node a which has at least two edges to nodes b and c. Having

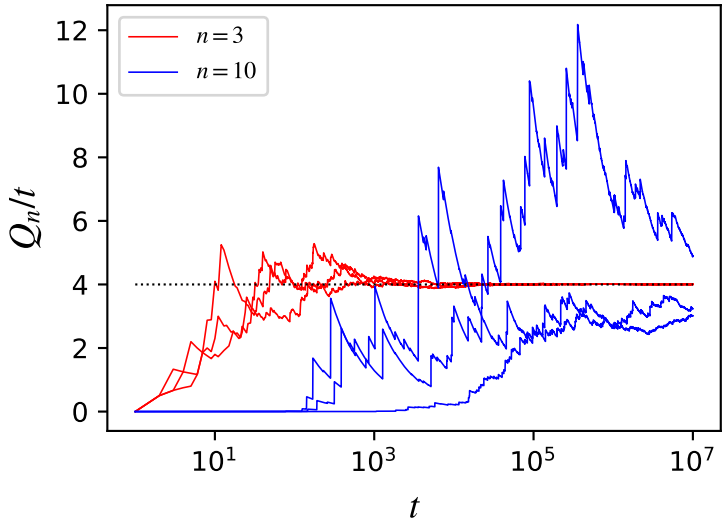

Figure 3.5: The rescaled number of cliques of size n, Q_n/t, as a function of time, t. Each red (blue) line corresponds to a single simulation tracking $n = 3$ ($n = 10$) cliques. The figure shows that after some small initial fluctuations, the number of cliques of size $n = 3$ quickly converges to our expected scaling. The number of cliques of size $n = 10$ shows significantly larger fluctuations and does not converge to our expected scaling by the end of the simulation duration, $t = 10^7$. However, all three traces do appear to slowly approach the $Q_n/t = 4$ scaling towards the end of the simulation.

identified this wedge, what is the probability that an edge exists between nodes b and c, forming a triangle? In the case of a random graph the probability of closing the wedge goes to zero as $t \to \infty$. However, in the CCM, where triangle closure is deterministic, it is guaranteed that for any set of three nodes connected by two edges, the third edge closing the wedge is occupied with probability 1. Hence, a triangle is formed in the CCM but not in the random graph, resulting in graphs with the same average edge occupation probability but different clique structure. A similar argument can be made for the deterministic formation of larger cliques in the CCM. Considering differences in cliques scaling across a wider range of growing network models would be a valuable exercise for the future.

3.3.4 Clustering

Transitivity is a global clustering measure defined as

$$\tau_{G_O} = 3 \times \frac{\#(\text{triangles in } G_O)}{\#(\text{twigs in } G_O)}, \tag{3.33}$$

where a twig is any three nodes connected by two edges. The number of twigs is equivalent to the number of star graphs of size 2, s_2, where a star graph of size n is a subgraph with 1 central node and n connected neighbours. The number of subgraphs of size 2 is related to the degree distribution by

$$s_2(t) = t \sum_{k_O \geq 2} \binom{k_O}{2} \cdot p_O(k_O) = t \cdot \frac{\langle k_O^2 \rangle - \langle k_O \rangle}{2}, \tag{3.34}$$

where we have used the property that $p_O(k < 2) = 0$. Recalling that $\langle k_O \rangle = 6$ and $\langle k_O^2 \rangle \approx 62$, the number of twigs scales as $s_2 \sim 28t$, such that

$$\tau_{G_O} = \frac{3Q_3}{s_2} \sim \frac{3 \cdot 4t}{28t} = \frac{3}{7}. \tag{3.35}$$

Note that the observed network can be recovered from the hidden network by converting every wedge in G_H into a triangle. This can be thought of as complete triadic closure where every possible triangle which can be closed, from the addition of a single edge to the hidden network, is closed. This implies that the CCM has the largest possible transitivity from a single iteration of triadic closure on a random recursive tree.

The local clustering coefficient, $cc(\alpha)$, is defined as the number of edges between the $(k_O)_\alpha$ neighbours of α, normalised by the the number of edges in a complete subgraph of size $(k_O)_\alpha$. For the CCM,

$$cc(\alpha) = \frac{\binom{(k_H)_\alpha}{2} + \sum_{\beta=1}^{(k_H)_\alpha} \binom{(k_H)_{\alpha,\beta}}{2}}{\binom{(k_O)_\alpha}{2}}, \tag{3.36}$$

where the first term corresponds to the complete subgraph of the $(k_H)_\alpha$ neighbours of α in G_H, and the sum contributes the edges from one complete subgraph formed by node α, β and its $(k_H)_{\alpha,\beta} - 1$ neighbours, excluding α. The global clustering coefficient, $CC(G_O)$, is defined as the average of Eq. (3.36) over all nodes in the network. In simulations, $CC(G_O) \approx 0.771$ for large t.

3.3.5 Path lengths

Steinbock, Biham, and Katzav [108] calculate the distribution of shortest path lengths for the UCM (referred to in their paper as the corded node duplication model). Specifically, the authors calculate the probability that two randomly chosen nodes, i and j, will be separated by a shortest path of length ℓ, denoted as $\mathcal{P}(L = \ell; t)$, at time t.

The UCM with $p = 0$ corresponds to a random recursive tree and is therefore equivalent to the G_H. Hence, for the hidden network, we can lift the path length distribution, $\mathcal{P}_H(L_H = \ell; t)$, and the mean shortest path, $\langle L_H(t)\rangle$, from Steinbock, Biham, and Katzav [108]. We can then exploit a convenient mapping to calculate the distribution of shortest path lengths in G_O from G_H.

Consider two randomly chosen nodes i and j. In G_H, there is a unique path (due to its tree structure) from i to j of length $(\ell_H)_{ij}$. In G_O, the enforced triadic closure process means that for every two steps on the path from i to j in G_H, an observed edge exists in G_O which acts as a shortcut, reducing the path length by one. Hence, if the path length $(\ell_H)_{ij}$ is even, the path length in G_O is given by $\ell_O = \ell_H/2$; if the path length is odd, $\ell_O = (\ell_H + 1)/2$. Using this mapping, we can write

$$\mathcal{P}_O(L_O = \ell; t) = \mathcal{P}_H(L_H = 2\ell; t) + \mathcal{P}_H(L_H = 2\ell - 1; t), \text{for } \ell \geq 1. \tag{3.37}$$

If we assume that, for large t, there are an approximately equal number of odd and even shortest paths in G_H, the average shortest path length in G_O is

$$\langle L_O(t)\rangle = \frac{\langle L_H(t)\rangle}{2} + \frac{1}{4}, \tag{3.38}$$

where the $1/4$ term accounts for the discrepancy in the mapping for odd and even paths.

From Steinbock, Biham, and Katzav [108], we note that the mean shortest path length for G_H scales as

$$\langle L_H(t)\rangle \sim 2 \cdot \ln(t), \tag{3.39}$$

which indicates that the hidden network exhibits the small-world property [11]. We have omitted constants which are negligible at large t. Hence, applying the mapping in Eq. (3.38) and omitting the 1/4 term for simplicity, the mean shortest path length for G_O is given by

$$\langle L_O(t) \rangle \sim \ln(t), \tag{3.40}$$

indicating that the observed network also exhibits the small-world phenomenon. This mapping is confirmed by simulations.

For interest, we note that for $0 < p < 1$, the shortest paths for the UCM are in general not unique; there may be multiple paths between nodes i and j which are equally short. Unusually for a non-tree network, all shortest paths are unique in the CCM.

3.4 General correlated copying model

The general correlated copying model (GCCM), is defined analogously to the CCM, starting with observed and hidden networks initialised at $t = 1$. Like the UCM and CCM, the GCCM is a corded node duplication model. For practical reasons, we initialise the graph with three nodes which form a complete graph in G_O, and a wedge in G_H. This ensures that the initial graph contains some edges found in G_H, and some edges found only in G_O.

At $t = t_\alpha$, node α is added to both networks and a single target node, β, is chosen uniformly at random. We label the k_O^β neighbours of β in G_O with the index γ_j. In the observed network, the copied edge (α, γ_j) is formed with probability p_{hid} if the edge $(\beta, \gamma_j) \in E_H$ (inner circle copying), and probability p_{obs} otherwise (outer circle copying). The direct edge (α, β) is added to both G_O and G_H.

The GCCM encapsulates a wide spectrum of heterogeneous copying. Setting $p_{\text{hid}} = 1$ and $p_{\text{obs}} = 0$ reduces the GCCM to the CCM, whereas setting $p_{\text{hid}} = p_{\text{obs}} = p$ reduces the GCCM to the UCM. We have discussed the social motivation for the case where $p_{\text{hid}} > p_{\text{obs}}$, representing a copying bias towards the inner social circle of a node. However, the GCCM can also be tuned

to the reverse case where $p_{\text{hid}} < p_{\text{obs}}$, resulting in a bias against inner circle nodes. We are not aware of a clear physical motivation for this latter case. However, the structural diversity of these anti-correlated networks warrants their discussion here.

Figure 3.6 shows numerical results for (a) the effective copying probability, (b) the densification exponent, which describes the growth in time of the number of edges in the observed network, (c) the average local clustering coefficient, and (d) the transitivity, for the GCCM with 10^4 nodes.

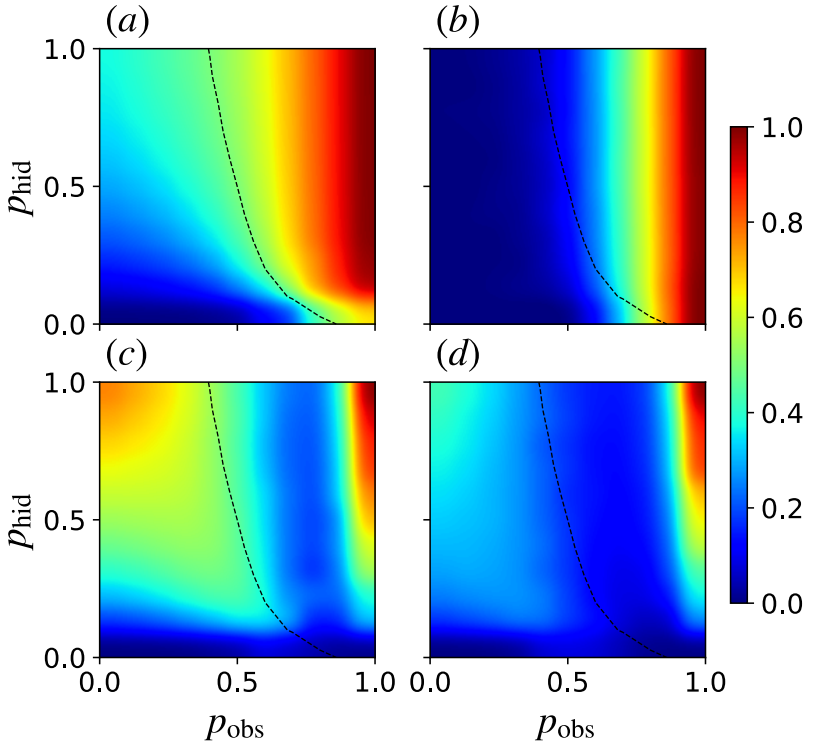

Figure 3.6: Properties of the observed network, G_O, in the general correlated copying model (GCCM). Numerical results for 10^4 nodes as a function of the hidden copying probability, p_{hid}, and the outer copying probability, p_{obs}. (a) The effective copying probability. (b) The densification exponent. (c) The average local clustering coefficient, $CC(G_O)$. (d) The transitivity, τ_{G_O}. Black dashed contour: effective copying probability of 0.5 at $t = 10^4$, calculated numerically. Values have been smoothed for clarity.

The effective copying probability corresponds to the fraction of target node neighbours which appear to be copied in the observed network. Formally, we can write the average effective

copying probability at time t as

$$p_{\text{eff}}(t) = \left\langle \frac{p_{\text{hid}} k_H^\beta + p_{\text{obs}}(k_O^\beta - k_H^\beta)}{k_O^\beta} \right\rangle, \tag{3.41}$$

where β is the index of the target node at time t, the first term represents edges copied from node β's inner circle, and the second term represents edges copied from the outer circle.

The dashed contour in Fig. 3.6(b) corresponds to an effective copying probability of 0.5, calculated numerically by averaging over the preceding 10^4 time steps. We note that $p_{\text{eff}} = 0$ if $p_{\text{hid}} = p_{\text{obs}} = 0$ (random tree), $p_{\text{eff}} = 1$ if $p_{\text{hid}} = p_{\text{obs}} = 1$ (complete graph), and $p_{\text{eff}} = p$ if $p_{\text{hid}} = p_{\text{obs}} = p$ (UCM). In general, the rise in p_{eff} is faster with increasing p_{obs} than increasing p_{hid}, although for $p_{\text{hid}} = 0$ we find very small p_{eff}, even for large p_{obs}. However, this observation is somewhat deceptive since, if the GCCM is in the dense regime and $p_{\text{hid}} \neq p_{\text{obs}}$, p_{eff} is not stationary. Calculated over longer time frames, we note that the effective copying probability appears to slowly converge to the outer copying probability, $p_{\text{eff}} \to p_{\text{obs}}$, since for $t \to \infty$, the ratio of the number of edges in the hidden network to the number of edges in the observed network tends to zero. This suggests that the dashed $p_{\text{eff}} = 0.5$ contour will converge to the $p_{\text{obs}} = 0.5$ line as $t \to \infty$.

We test whether the GCCM is in the sparse or dense regime explicitly by tracking the growth in the number of edges in the observed network. Let us define the densification exponent, ε, $0 \leq \varepsilon \leq 1$, using $E_O(t) \propto t^{1+\varepsilon}$, which relates the number of edges in the observed network to the number of nodes t. If $\varepsilon \approx 0$, the GCCM is sparse. If $\varepsilon = 1$, the GCCM grows as a complete graph. For intermediate values, the GCCM undergoes densification. For the UCM, the transition from the sparse to dense regime is known to take place at $p = 0.5$ [62]. We have not analytically calculated the transition for the GCCM, but may intuitively expect the transition at $p_{\text{obs}} = 0.5$ since the hidden network is a random tree. This seems to be supported by the numerical values of ε in Fig. 3.6(b), although the transition from zero to non-zero ε is shifted to slightly larger p_{obs} for $p_{\text{hid}} = 0$, and to smaller p_{obs} for $p_{\text{hid}} = 1$; this shift is likely to disappear as $t \to \infty$.

Figures 3.6(c) and (d) show the average local clustering coefficient, $CC(G_O)$, and transitivity (global clustering), τ_{G_O}, for the GCCM. Patterns are similar between the two figures, although local clustering generally exceeds global clustering in the sparse regime. For the UCM it is known that, in the dense regime, τ_{G_O} slowly converges to zero as $t \to \infty$, unless $p = 1$ [96]. In contrast, the local clustering appears to remain non-zero.

As expected, clustering is minimised at $p_{\text{hid}} = p_{\text{obs}} = 0$ (random tree) and maximised for a complete graph, $p_{\text{hid}} = p_{\text{obs}} = 1$. However, in the sparse regime we find that the maximum clustering is found at $p_{\text{hid}} = 1$, $p_{\text{obs}} = 0$ which corresponds to the CCM. Bhat *et al.* [96] note that local and global clustering for the UCM is not a monotonically increasing function of the copying probability p, with a local maxima in the sparse regime at non-zero p. This bimodal clustering is also present in the GCCM. In the anti-correlated regime where $p_{\text{hid}} \approx 0$, we find near zero clustering values. In particular if $p_{\text{hid}} = 0$ and $p_{\text{obs}} = 1$, we observe the unusual property that $\varepsilon \approx 1$, such that the network scales as (but is not) a complete graph, yet both the local and global clustering are approximately zero.

Extracting the degree distributions for the GCCM for various p_{hid} and p_{obs} shows similarly diverse behaviour, see Fig. 3.7. Each distribution is averaged over 100 realisations, but points are left deliberately unbinned to illustrate the significant fluctuations observed in the dense regime. For $p_{\text{hid}} = p_{\text{obs}} = 0$ (bottom left) the GCCM reduces to a random recursive tree, see Eq. (3.5). The CCM case with $p_{\text{hid}} = 1$, $p_{\text{obs}} = 0$ (top left) follows Eq. (3.16), where the tail can be approximated as a stretched exponential. This distribution is also shown in Fig. 3.2. Along the diagonal where $p_{\text{hid}} = p_{\text{obs}}$ (UCM), the degree distribution has a power law tail in the sparse regime, and exhibits anomalous scaling in the dense regime ($p \geq 0.5$). For $p_{\text{hid}} = p_{\text{obs}} = 1$, the GCCM reduces to a complete graph and all nodes have degree $t - 1$.

For $p_{\text{obs}} = 0$, the power law scaling observed in the UCM is completely suppressed, with a gradual transition from exponential decay to a stretched exponential tail as p_{hid} is increased from 0 to 1. In the sparse regime with $p_{\text{obs}} \neq 0$, all degree distributions appear fat tailed with only small deviations from the power laws observed for the UCM. However, unusual scaling is observed for $p_{\text{hid}} = 0$, $p_{\text{obs}} \neq 0$, where the distributions exhibit initial exponential decay

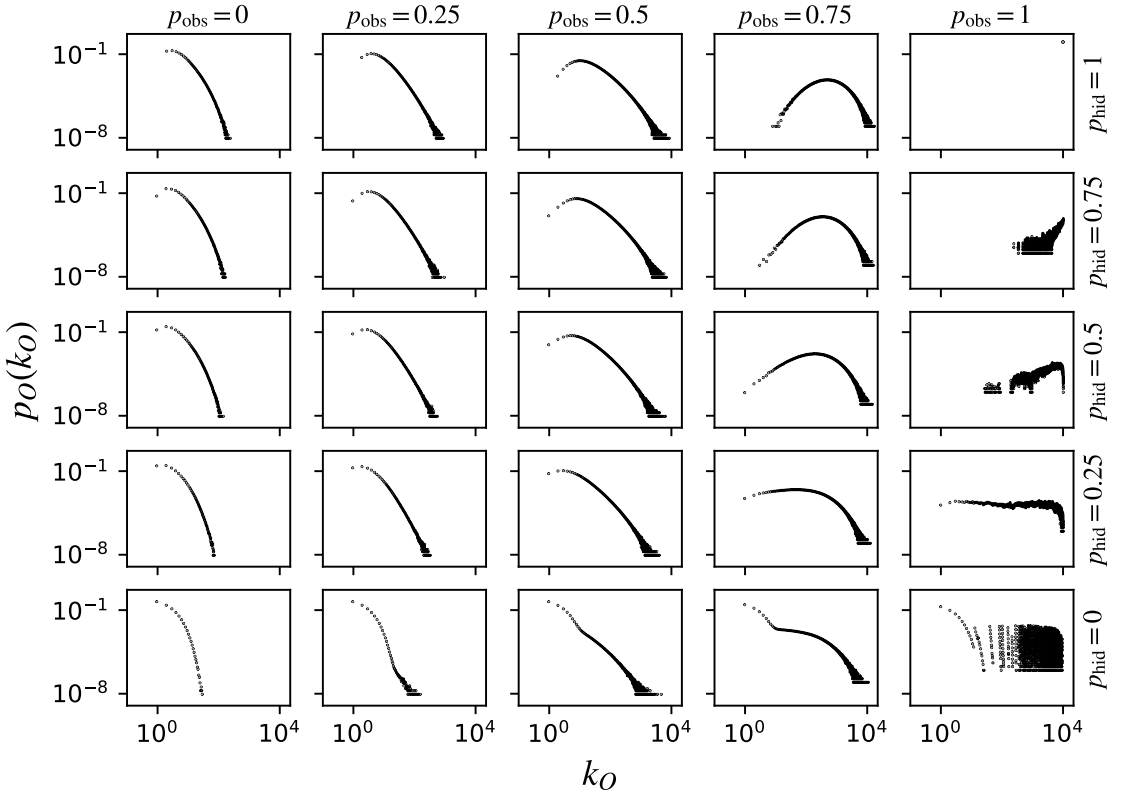

Figure 3.7: The observed degree distributions for the general correlated copying model (GCCM). The degree probability, $p_O(k_O)$, is plotted as a function of the observed degree, k_O, for various values of the outer copying probability, p_{obs} (left to right), and the hidden copying probability, p_{hid} (bottom to top). For $p_{\text{obs}} \in \{0, 0.25, 0.5\}$, each network contains 10^6 nodes. For $p_{\text{obs}} = 0.75$, each network contains 10^5 nodes. For $p_{\text{obs}} = 1$, each network contains 10^4 nodes. Distributions are averaged over 100 realisations. In the dense regime ($p_{\text{obs}} > 0.5$), network growth is non-ergodic leading to anomalous scaling and noisy degree distributions. The distribution at $p_{\text{hid}} = p_{\text{obs}} = 0$ corresponds to a random recursive tree, see Eq. 3.5 (exponential decay). The distribution at $p_{\text{hid}} = p_{\text{obs}} = 1$ corresponds to a complete graph. The distribution at $p_{\text{hid}} = 1$, $p_{\text{obs}} = 0$ corresponds to the CCM, see Fig. 3.2. If $p_{\text{hid}} = p_{\text{obs}}$, the GCCM is equivalent to the UCM.

at small k_O, attributable to the hidden network, before a second fat-tailed regime starting at intermediate k_O.

In the dense regime, all degree distributions exhibit anomalous scaling, such that individual realisations are not self-averaging. For $p_{\text{obs}} = 0.75$, the tail of the degree distributions is largely consistent across all p_{hid}. However, the probability of finding nodes with small degree is large for $p_{\text{hid}} = 0$, and is gradually suppressed as $p_{\text{hid}} \to 1$. These effects are most pronounced for $p_{\text{obs}} = 1$ where the modal degree is 1 for $p_{\text{hid}} = 0$, and $t-1$ for $p_{\text{hid}} = 1$, with a gradual transition

in between. Throughout this transition, the degree distribution appears almost uniform at $p_{\text{hid}} = 0.25$, where the probability of finding nodes with any given degree is approximately constant up until the large k_O limit. However, this effect is only observed when averaging over many realisations, with a much smaller degree range observed in individual networks.

3.5 Comparing copying models

We have introduced a simple model of heterogeneous node copying, motivated by arguments that triadic closure may not be structurally homogeneous in real networks.

Comparing the CCM, for which we have analytical results, to the UCM with the equivalent effective copying probability ($p = p_{\text{eff}} = 0.374$) we find significant differences in network structure. Both the average local clustering coefficient, $CC(G_O)$, and the transitivity, τ_{G_O}, are significantly larger in the CCM than the UCM. The CCM suppresses the power law tail observed in the UCM for the sparse regime, and consequently, the degree variance observed in the CCM is smaller than for the UCM. CCM: $\sigma^2(k_O) \approx 26$; UCM: $\sigma^2(k_O) \approx 192$. The CCM also has the unusual property, not found in the UCM, that the growth in the number of cliques of size n scales independently of n as $t \to \infty$. For both the UCM and CCM, the mean shortest path lengths scale as $\sim \ln(t)$ indicative of the small-world property.

The above comparison uses a single effective copying probability, but key differences are robust for variable p in the sparse regime. Specifically, the UCM degree distribution always exhibits a power law tail [62], and the largest measured clustering coefficients fall below the values seen for the CCM, see Tab. 3.1. Relaxing the CCM to the GCCM, we note that for large p_{hid} and small p_{obs}, the measured clustering values regularly exceed those observed in the UCM, with the UCM only reaching similar values far into the dense regime. Given the continuing debate about the ubiquity of power laws in real networks [19], the observation that power laws are suppressed in the GCCM when the UCM symmetry is broken supports the view that power law network scaling is an idealised case which in practice is rarely observed for real networks.

Whether such extreme bias is plausible in real networks is uncertain. However, observations

	p_{eff}	$CC(G_O)$	τ_{G_O}	**Degree Distribution**
CCM	0.37	0.77	0.43	Stretched Exponential
UCM (Max $CC(G_O)$)	0.38	0.52	0.20	Power law
UCM (Max τ_{G_O})	0.22	0.40	0.28	Power law

Table 3.1: A comparison between the correlated copying model (CCM) and the uniform copying model (UCM). UCM is simulated twice, once with an effective copying probability, p_{eff}, that results in the network with the highest average local clustering coefficient (in the sparse regime), $CC(G_O)$, and once with the effective copying probability that gives the largest transitivity, τ_{G_O}. Values are averaged over 50 simulations where each network contains 10^5 nodes. Standard deviations are negligible. All values for the observed network, G_O.

in academic collaboration networks suggest that extreme bias may be possible [124]. For instance, Kim and Diesner [124] show that the ratio of triadic closure between two nodes is approximately zero if the number of shared collaborators is zero, rises rapidly as the number of shared collaborators increases, and plateaus at a ratio of one.

A second clue towards heterogeneous copying is the observation of very large clustering values in real networks. A selection of these networks and their clustering coefficients is shown in Tab. 3.2. Stressing that both the UCM and GCCM are toy models of node copying, the networks in Tab. 3.2 exhibit average local and/or global clustering far exceeding even the most optimistic values for the UCM, likely reflecting the clique structure of the underlying networks. In contrast, the listed clustering values are relatively similar to what may plausibly emerge from heterogeneous copying, although even the clustering observed for the extreme CCM case falls below some of the values shown in Tab. 3.2. Future work should go beyond this qualitative analysis and should attempt to measure the degree to which copying symmetry is broken for real networks where these mechanisms are relevant.

3.6 Discussion

The UCM, CCM and GCCM are all examples of corded copying models where an edge forms between a newly added node and the target node which is duplicated. This is in contrast to uncorded duplication models where a new node is formed by copying an existing target node

	Nodes	**Edges**	$CC(G_O)$	τ_{G_O}
arXiv Astro coauthors	18.8K	198.1K	0.63	0.32
arXiv GR coauthors	5.2K	14.5K	0.53	0.63
arXiv CM coauthors	23.1K	93.4K	0.63	0.26
arXiv HEP coauthors	22.9K	2.7M	0.81	0.31
DBLP coauthors	540.5K	15.2M	0.80	0.65
NetSci coauthors	379	914	0.74	0.43
Hollywood collaborations	1.1M	56.3M	0.77	0.31
DNC Email corecipients	906	12.1K	0.61	0.56

Table 3.2: A selection of sparse undirected networks which may plausibly grow via a copying mechanism. These networks exhibit larger average local clustering coefficients, $CC(G_O)$, and transitivity, τ_{G_O}, than one may expect if these networks were to grow via a uniform copying mechanism. Network source data from Rossi and Ahmed [130].

and its neighbours, but an edge is not formed between the new node and the duplicated target node. Corded models are more common in the context of social phenomena and triadic closure, whereas uncorded models are typically more relevant to duplication–divergence processes in protein interaction networks. In the current work, we have focused exclusively on corded models; considering heterogeneous copying in uncorded models [105, 106, 114, 115] would be an appropriate future extension. Heterogeneous copying could also be studied by extending directed copying models [116, 117].

The GCCM and CCM are examples of hidden network models. From a mathematical standpoint, hidden network models can be thought of as a variant of interdependent networks where nodes in one layer have dependencies of nodes in another layer [16, 128]. However, at a conceptual level, hidden network models puts an emphasis on how the evolution of network structure can depend on asymmetries not observed in our data.

In this chapter, we have focused on copying in social networks, but the ideas naturally extend to other contexts. In economics, our framework may be applied to shareholder networks [131], where nodes are connected if they both own a common asset. Here, the hidden network represents the full set of co-owned assets, whereas the observed network includes publicly disclosed assets. Similarly, the idea can be applied to co-bidding networks in public procurement,

where an edge indicates that two companies both placed bids on the same contract. In many jurisdictions, only winning bids (of which there may be multiple) are publicly revealed. Therefore, the observed network may represent the network of winning bids, whereas the hidden network includes all bids. Hidden network models may be a valuable representation in these cases if there are structural reasons for why some data is observed and some data is hidden. For instance, fraudulent behaviour in public procurement has been associated with anomalous structural features in the co-bidding network [132].

Other examples may be found in ecology, where multilayer networks have been used to represent different interactions between a common set of species [133, 134]. Kéfi *et al.* [133] find that the structure of interactions in one layer has significant cross-dependencies to the structure of other layers. This mirrors how interlayer dependencies in the CCM are used to break symmetries in the evolution of the observed network. Finally, hidden networks may find general relevance to other fields where interdependent networks have been influential. This may include studies on energy demand management for power grids [135], and the emergence of synchronisation in multilayer neuronal models [136].

A more unusual application of the hidden network concept is for decomposing complex single layer networks into simpler two-layer structures. One such example is second-neighbour preferential attachment, the k2 model (see chapter 2), an implementation of the Barabási-Albert model where nodes attach proportionally to the number of nodes within two steps of a target node [1]. Using our framework, the model is decomposed into an observed network, and a hidden network (in this case referred to as the influence network) where nodes are connected to all nodes which are two or fewer steps away, representing the node's sphere of influence. Here, second-neighbour preferential attachment is equivalent to conventional first-neighbour preferential attachment followed by a local copying step. Structural heterogeneity that is intrinsic in such a model has profound consequences for the time dependence of network growth [1].

Overall, we have introduced a general model of heterogeneous copying, implemented using the hidden network model framework. In the case of extreme copying bias, we have derived analytical results and have demonstrated significant differences to similar models with uniform

copying rules. In particular, power-law degree distributions observed in uniform copying can be suppressed under heterogeneous copying, and networks are significantly more clustered if copying is biased towards a node's inner circle. Although a systematic study of copying in real networks is necessary, evidence suggests that heterogeneous copying may be relevant in a social context.

3.7 Conclusions

The previous two chapters have focused on two different models of growing networks: In chapter 2, I discussed the k2 model whose purpose is to illustrate the role of time dependence in network growth, particularly with regards to non-stationary growth rules. Such non-stationarity is observed for a range of real-world networks on long time scales. However, I have not discussed in depth the precise nature of the non-stationarity for these networks. In chapter 3, I introduced the correlated copying model (and its general extension) to investigate the role of structural heterogeneity in copying models. Although largely stand-alone, together these models reflect my interest in probing how key assumptions in the construction of theoretical network models affects our beliefs about real world networks. In the previous two chapters, we have not studied real networks in detail. However, simple observations from real data clearly indicate that more work is needed to fully understand the implication of these assumptions.

As I noted in this part's overview, a common theme which runs through the previous two chapters is the use of hidden network models which offer both a valuable conceptual framework for thinking about network models, and a powerful mathematical tool for analysing them. In this context, the idea which interests me most regards the role of the causal coupling between the observed and the hidden network layers (referred to as the influence network in chapter 2). Consider, for instance, the correlated copying model introduced in the current chapter. Here, the structure of the hidden network determines which edges are copied in the observed network. Therefore, if we perturb the edge structure of the hidden network, this will have a causal effect on the influence of the observed network. However, the reverse is not true – perturbing the

observed network has no causal effect on the hidden network. It is this asymmetry which, in part, is responsible for our ability to analyse the correlated copying model analytically.

What is interesting for me, is that this isn't the case in the k2 model, where causal coupling between the observed and hidden (influence) layers is bidirectional, and consequently, the k2 model has evaded any of my attempts at a robust analytical treatment. The correlated copying model and the k2 model are specific examples of hidden network models with uni- and bidirectional coupling respectively. Although the consequences of different forms of coupling is not fully clear, my aim is to explore this problem in the future.

3.8 Chapter summary

- I introduce the correlated copying model (CCM) to investigate heterogeneity in node copying. The CCM is studied both analytically and numerically.
- The CCM is constructed using hidden network models, a framework which decomposes a complex network model into an observed and hidden layer.
- The CCM demonstrates significant differences to the equivalent model with homogeneous copying rules, the uniform copying model (UCM). The CCM suppresses power law degree distributions observed in the UCM, and is significantly more clustered. In fact, the CCM is the network with the highest possible transitivity that is constructed from a single round of triadic closure on a random recursive tree.
- The CCM has the unusual property, proven analytically, that the number of cliques of size n grows independent of n.
- I introduce a general form of the correlated copying model (the GCCM), which interpolates between the CCM and the UCM. The GCCM is studied numerically.
- With reference to a selection of real-world networks, I suggest that heterogeneous copying mechanisms may be of particular relevance to social networks, particularly collaboration networks.

3.9 Mathematical derivations

3.9.1 Observed network second moment

Here we calculate the second moment for the observed network, $\langle k_O^2 \rangle$. First, note that we can write the observed degree of node α as

$$(k_O)_\alpha = \sum_{\beta=1}^{(k_H)_\alpha} (k_H)_{\alpha,\beta}. \tag{3.42}$$

Let us square both sides, and average over all nodes in the network such that

$$\frac{1}{t}\sum_{\alpha=1}^{t}(k_O)_\alpha^2 = \frac{1}{t}\sum_{\alpha=1}^{t}\left[\sum_{\beta=1}^{(k_H)_\alpha}(k_H)_{\alpha,\beta}\right]^2. \tag{3.43}$$

We now simplify the left hand side and expand the squared sum on the right hand side to give

$$\langle k_O^2 \rangle = \frac{1}{t}\sum_{\alpha=1}^{t}\left[\sum_{\beta=1}^{(k_H)_\alpha}(k_H)_{\alpha,\beta}^2 + \sum_{\beta=1}^{(k_H)_\alpha}\sum_{\beta'\neq\beta}^{(k_H)_\alpha}(k_H)_{\alpha,\beta}(k_H)_{\alpha,\beta'}\right], \tag{3.44}$$

where the first term in the brackets is the contribution from the product of node α, β with itself, and the second term considers the cross term contribution between node α, β and all other neighbours α, β'. Taking the first term out of the brackets and summing over all nodes in the network gives

$$\langle k_O^2 \rangle = \langle k_H^3 \rangle + \frac{1}{t}\sum_{\alpha=1}^{t}\sum_{\beta=1}^{(k_H)_\alpha}\sum_{\beta'\neq\beta}^{(k_H)_\alpha}(k_H)_{\alpha,\beta}(k_H)_{\alpha,\beta'}, \tag{3.45}$$

where we have used the property that node ℓ appears in the expanded sum $(k_H)_\ell$ times. We proceed with a mean-field approximation and assume that as $t \to \infty$ the average degree of next-nearest neighbours in the hidden network are uncorrelated. This implies that $(k_H)_{\alpha,\beta}$ and $(k_H)_{\alpha,\beta'}$ are uncorrelated. Hence, we can replace the sum over β' with $(k_H)_\alpha - 1$ times the

average degree of node $(k_H)_{\alpha,\beta'}$ giving

$$\langle k_O^2 \rangle \approx \langle k_H^3 \rangle + \frac{1}{t} \sum_{\alpha=1}^{t} \sum_{\beta=1}^{(k_H)_\alpha} (k_H)_{\alpha,\beta} \langle k_H \rangle_{\alpha,\beta'} \cdot ((k_H)_\alpha - 1) \,. \tag{3.46}$$

We know that due to the friendship paradox $\langle k_H \rangle_{\alpha,\beta'} \neq \langle k_H \rangle$. On average our friends have more friends than we do. For a network with zero assortativity, we can relate the average nearest neighbour degree to the average degree by

$$\langle k_H \rangle_{\alpha,\beta'} = \langle k_H \rangle + \frac{\sigma^2(k_H)}{\langle k_H \rangle}, \tag{3.47}$$

where σ^2 is the degree variance. In our case this gives

$$\langle k_H \rangle_{\alpha,\beta'} = \langle k_H \rangle + \frac{\langle k_H^2 \rangle - \langle k_H \rangle^2}{\langle k_H \rangle} = 3. \tag{3.48}$$

Substituting this into Eq. (3.46) and rearranging slightly we can write

$$\langle k_O^2 \rangle \approx \langle k_H^3 \rangle + \frac{3}{t} \sum_{\alpha=1}^{t} \sum_{\beta=1}^{(k_H)_\alpha} \left[(k_H)_{\alpha,\beta} (k_H)_\alpha - (k_H)_{\alpha,\beta} \right], \tag{3.49}$$

where again we have a cross term between node α, β and node α. Since the inner sum does not act on node α, we may also assume that the $(k_H)_\alpha$ term can be replaced by Eq. (3.46) such that

$$\langle k_O^2 \rangle \approx \langle k_H^3 \rangle + \frac{3}{t} \sum_{\alpha=1}^{t} \sum_{\beta=1}^{(k_H)_\alpha} \left[3(k_H)_{\alpha,\beta} - (k_H)_{\alpha,\beta} \right] = \langle k_H^3 \rangle + \frac{6}{t} \sum_{\alpha=1}^{t} \sum_{\beta=1}^{(k_H)_\alpha} (k_H)_{\alpha,\beta}. \tag{3.50}$$

Finally, noting that each $(k_H)_{\alpha,\beta}$ term appears exactly $(k_H)_{\alpha,\beta}$ times in the expanded sum we find

$$\langle k_O^2 \rangle \approx \langle k_H^3 \rangle + 6 \langle k_H^2 \rangle = 62, \tag{3.51}$$

where in the final line we have used $\langle k_H^3 \rangle = 26$ and $\langle k_H^2 \rangle = 6$, both of which follow directly from the hidden degree distribution, $p_H(k_H) = 2^{-k_H}$. This result is consistent with direct measurements taken in simulations.

3.9.2 Observed degree distribution

For the observed network we start with the master equation

$$N(t+1)\cdot p_O(k_O,t+1) - N(t)\cdot p_O(k_O,t) =$$
$$N(t)\cdot \Pi_O(k_O-1,t)\cdot p_O(k_O-1,t) - N(t)\cdot \Pi_O(k_O,t)\cdot p_O(k_O,t) + 2^{1-k_O},\ \text{for } k_O \geq 2, \quad (3.52)$$

where $N(t)$ is the number of nodes at time t, $\Pi_O(k_O,t)$ is the probability that a node with observed degree k_O gains a new edge, and $p_O(k_O,t)$ is the degree distribution. The first line corresponds to nodes whose observed degree remains unchanged from t to $t+1$, the first and second terms on the second line correspond to the transitions $k_O - 1 \to k_O$ and $k_O \to k_O + 1$ respectively, and the final term corresponds to the probability that the newly added node has initial degree k_O, given by $p_H(k_O - 1)$. Letting $N(t) = t$ and assuming the degree distribution is stationary for large t,

$$p_O(k_O) = \pi_O(k_O-1)\cdot p_O(k_O-1) - \pi_O(k_O)\cdot p_O(k_O) + 2^{1-k_O},\ \text{for } k_O \geq 2, \quad (3.53)$$

where

$$\pi_O(k_O) = t\cdot \Pi_O(k_O) = t\cdot \frac{1+\langle k_H \mid k_O\rangle}{t} = 1 + \langle k_H \mid k_O\rangle, \quad (3.54)$$

and $\langle k_H \mid k_O\rangle$ is the average degree of nodes in the hidden network with a fixed observed degree of k_O. Here, the $1/t$ term corresponds to the probability that edges are gained from direct attachment, whereas the $\langle k_H \mid k_O\rangle/t$ corresponds to the probability that edges are gained from copying.

3.9.3 Derivation of Eq. (3.9)

Consider the node α added to the network at time t_α. The initial conditions for node α are

$$\begin{aligned}(k_H(t_\alpha))_\alpha &= 1,\\ \langle k_O(t_\alpha)\rangle_\alpha &= 1 + \langle k_H(t_\alpha - 1)\rangle_\beta,\end{aligned} \tag{3.55}$$

where the final term is the average hidden degree of the target node β. In the hidden network, node α gains edges from direct attachment only. Hence, at $t > t_\alpha$,

$$\langle k_H(t > t_\alpha)\rangle_\alpha = 1 + \sum_{j=t_\alpha}^{t-1} \frac{1}{j} = 1 + H_{t-1} - H_{t_\alpha - 1}, \tag{3.56}$$

where H_n is the n^{th} harmonic number. In the observed network, either node α is targeted via direct attachment, or a copied edge is formed from the new node to node α via any of the $(k_H(t))_\alpha$ neighbours of node α. Hence,

$$\begin{aligned}\langle k_O(t > t_\alpha)\rangle_\alpha &= \langle k_O(t_\alpha)\rangle_\alpha + \sum_{j=t_\alpha}^{t-1} \frac{1 + \langle k_H(j)\rangle_\alpha}{j}\\ &= \langle k_O(t_\alpha)\rangle_\alpha + \sum_{j=t_\alpha}^{t-1} \frac{2 + H_j - H_{t_\alpha - 1} - 1/j}{j},\end{aligned} \tag{3.57}$$

where we have subbed in Eq. (3.56) and $H_{j-1} = H_j - 1/j$. Note that

$$\sum_{j=1}^{n} \frac{H_j}{j} = \frac{1}{2}\left[(H_n)^2 + H_n^{(2)}\right], \tag{3.58}$$

where $H_n^{(m)}$ is the n^{th} generalised Harmonic number of order m defined as

$$H_n^{(m)} = \sum_{j=t_\alpha}^{n} \frac{1}{j^m}. \tag{3.59}$$

Hence, we can rewrite Eq. (3.57) as

$$\langle k_O(t > t_\alpha)\rangle_\alpha = \langle k_O(t_\alpha)\rangle_\alpha + (2 - H_{t_\alpha - 1}) \sum_{j=t_\alpha}^{t-1} \frac{1}{j} + \sum_{j=t_\alpha}^{t-1} \frac{H_j}{j} - \sum_{j=t_\alpha}^{t-1} \frac{1}{j^2}, \tag{3.60}$$

which expanded gives

$$\begin{aligned}\langle k_O(t > t_\alpha)\rangle_\alpha = \langle k_O(t_\alpha)\rangle_\alpha + (2 - H_{t_\alpha - 1})(H_{t-1} - H_{t_\alpha - 1}) \\ + \frac{1}{2}\left[(H_{t-1})^2 + H^{(2)}_{t-1} - (H_{t_\alpha - 1})^2 - H^{(2)}_{t_\alpha - 1}\right] - (H^{(2)}_{t-1} - H^{(2)}_{t_\alpha - 1}).\end{aligned} \tag{3.61}$$

We can simplify Eq. (3.61) slightly

$$\begin{aligned}\langle k_O(t > t_\alpha)\rangle_\alpha = \langle k_O(t_\alpha)\rangle_\alpha + (2 - H_{t_\alpha - 1})(H_{t-1} - H_{t_\alpha - 1}) \\ + \frac{1}{2}\left[(H_{t-1} + H_{t_\alpha - 1})(H_{t-1} - H_{t_\alpha - 1}) - H^{(2)}_{t-1} + H^{(2)}_{t_\alpha - 1}\right],\end{aligned} \tag{3.62}$$

followed by

$$\langle k_O(t > t_\alpha)\rangle_\alpha = \langle k_O(t_\alpha)\rangle_\alpha + \frac{1}{2}\left[(4 + H_{t-1} - H_{t_\alpha - 1})(H_{t-1} - H_{t_\alpha - 1}) - H^{(2)}_{t-1} + H^{(2)}_{t_\alpha - 1}\right], \tag{3.63}$$

as required.

Part 2

Modelling atrial fibrillation

Chapter 4

Overview & motivation

The human heart is made up of four chambers, two small atria at the top, and two larger ventricles at the bottom, see Fig. 4.1. Collectively, these chambers are responsible for pumping blood around the body. The right atrium collects deoxygenated blood from the body and transfers it to the right ventricle from where it is pumped to the lungs to be oxygenated. After being oxygenated, blood is received by the left atrium, passes to the left ventricle, and is pumped around the body, before the whole cycle repeats.

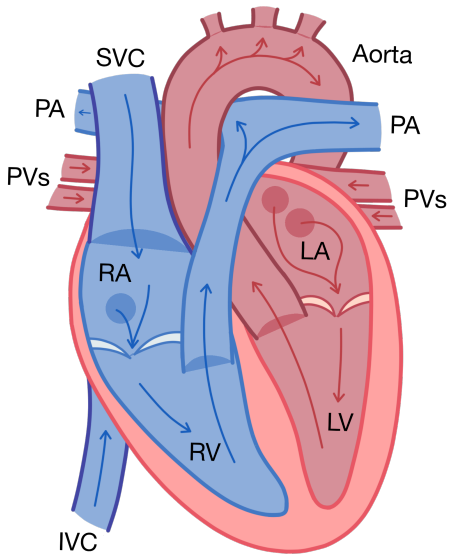

Figure 4.1: A schematic showing the four heart chambers, the right and left atria on top (RA/LA), and the right and left ventricles below (RV/LV). Arrows indicate the flow of oxygenated and deoxygenated blood in red and blue respectively. Key blood vessels labelled include the inferior and superior vena cava (IVC/SVC), the aorta, the pulmonary arteries (PAs), and the pulmonary veins (PVs). Schematic drawn by Becky Joakim.

In order to pump blood efficiently, the four heart chambers must contract in a precise and coordinated manner. To achieve this, the heart has a sophisticated system of electrical pathways which act as a signalling system, telling cardiac muscle cells when to contract. In the healthy heart, signals from the brain initiate an electrical wavefront at the sino-atrial node (SAN) in the right atrium, see Fig. 4.2. From here, the wavefront passes smoothly from the top of the atria to the bottom, through the branching structure of cylindrical heart muscle cells in which individual cells are well-connected along their longitudinal axis, but weakly transversely, mimicking a lattice like structure with variable coupling in the longitudinal and transverse directions. Electrical propagation through the network of muscle cells results in the smooth and regular contraction of the upper heart chambers. At the bottom of the atria, the wavefront is received by the atrio-ventricular node (AVN) which modulates the signal and sends an electrical wavefront through the ventricles, resulting in the coordinated contraction of the lower heart chambers. Note, that once a heart muscle cell has excited, it remains unexcitable for a short duration (the refractory period), such that each electrical wavefront is followed by an unexcitable wavetail. For a comprehensive review of electrical propagation in the heart see [137].

A cardiac arrhythmia forms when the smooth and regular propagation of electrical signal in the heart is disrupted, such that the resulting contraction of the heart muscle is no longer precisely coordinated. Cardiac arrhythmias come in many forms: some are largely benign with minimal impact on heart function, whereas others can result in sudden death within minutes; some have clear localised causes which are easily treated, whereas others are poorly understood and difficult to manage.

In the following two chapters, I investigate the mechanisms underlying atrial fibrillation (AF), an arrhythmia associated with the rapid and irregular beating of the heart's upper chambers. AF is the most common cardiac arrhythmia in clinical practice, yet, despite extensive research, there are significant disagreements within the cardiac electrophysiology community as to its underlying mechanisms [138–140]. The irregular and inefficient contraction of the atria during AF results in an increased risk of clotting, and, consequently, AF is a leading cause of stroke.

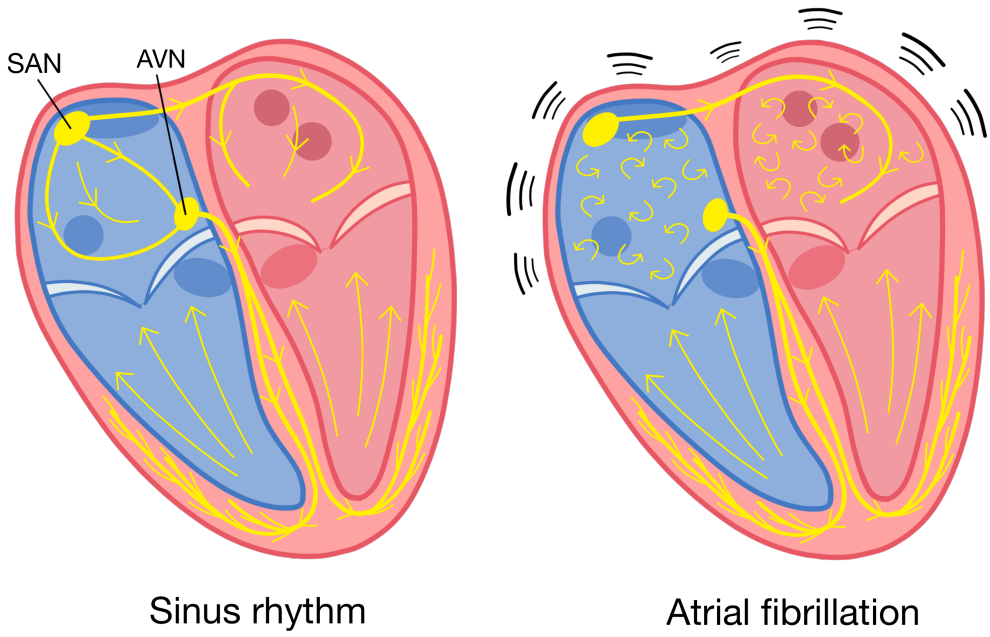

Figure 4.2: A schematic showing the heart's electrical system which coordinates the contraction of the four heart chambers. In the healthy heart (sinus rhythm), wavefronts propagate smoothly from the sino-atrial node (SAN) to the atrio-ventricular node (AVN), and from the right atrium to the left atrium. From the AVN, signals are passed along specialised fibres before wavefronts propagate up the ventricles, resulting in a bottom-to-top contraction. In atrial fibrillation, electrical conduction is disrupted in the atria with wavefronts propagating chaotically. Wavefronts remain planar in the ventricles due to the AVN's filtering mechanism, but are initiated at irregular intervals resulting in an irregular heart rhythm. Disrupted conduction in the atria leads to inefficient contraction with the atria appearing to quiver, resulting in a heightened risk of clotting and stroke. Schematic drawn by Becky Joakim.

Broadly speaking, AF is identified as the the chaotic propagation of electrical wavefronts in the atria. As such, understanding its mechanistic origin can be broken down into two parts: Firstly, how does AF initiate (i.e., how does the atria transition from a state where the electrical wavefronts propagate smoothly through the atria, to a state where wavefronts propagate chaotically), and secondly, once AF is initiated, how is chaotic propagation maintained? Most of the disagreements regarding the mechanisms of AF relate to this second question. In brief, one camp believes that AF is self-perpetuating in a delocalised and disorganised manner (often referred to a the multiple wavelet hypothesis), whereas the other camp believe that AF is perpetuated by local sources of fibrillatory (irregular and chaotically propagating) waves; both sides have accrued substantial experimental evidence in support of their views.

It is not my aim with this overview to discuss these arguments in detail; for more detail I happily

recommend the following reviews [138, 139]. However, I will note that, in my opinion, there is no need for the mechanisms of AF to be either organised, or disorganised; AF maintenance may well exhibit a spectrum of organisation which covers both cases, mediated by a range of factors related to the electromechanical conditions in the atria. Such a view has recently gained some experimental backing [141].

There are a range of managements for AF, both surgical and pharmacological. Amongst these, catheter ablation is generally considered the most effective [142]; a procedure where electrical catheters are passed into the atria via a patient's veins and used to burn and destroy regions on the inside of the heart which are thought to play a role in initiating or maintaining AF, see Fig. 4.3.

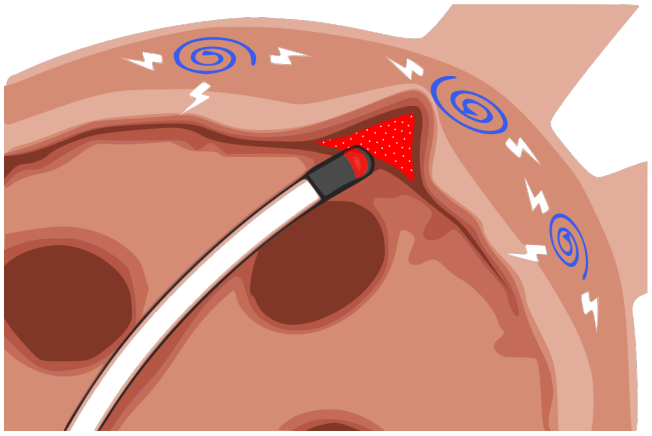

Figure 4.3: One approach to treating atrial fibrillation is to perform a catheter ablation, a procedure where a catheter is inserted into the atria and used to destroy or isolate the triggers or drivers of AF. If these drivers are localised (most often in the pulmonary veins), these drivers are easy to isolate or destroy. However, if there are too many drivers in the atria, it may not be feasible to destroy all of them without causing excessive atrial damage. Schematic drawn by Madalina Sas.

The current gold standard ablation technique is pulmonary vein isolation (PVI), a procedure aimed at isolating the electrical triggers which initiate AF. The procedure, which was developed in the late 1990s, followed a pioneering study which found that a large number of the electrical triggers which initiate AF originate in the muscular sleeves which extend into the pulmonary veins from the right atrium [143]. PVI aims to electrically disconnect these triggers from the atria, thereby preventing fibrillatory waves from the veins entering the atrial bulk. During PVI, a catheter is used to ablate the tissue that connects the left atrium to the veins which deliver

blood from the lungs. In doing so, the pulmonary veins become electrically disconnected from the atria. However, this treatment is only effective in around half of patients with AF, and often only for a few months or years before a patient experiences a relapse. Since its introduction, many variants of PVI ablation have been attempted, including the targeted ablation of focal points in the atrial bulk, but none of these methods have demonstrated a significant improvement over PVI. One of the motivations in this work is to consider why this might be the case.

In the following chapters, I focus on just one possible mechanism of AF initiation and maintenance, the formation of micro-anatomical reentrant circuits; continuously activate electrical circuits anchored to the micro-structure of the atria [144]. Loosely speaking, these can be thought of as microscopic electrical short circuits which form when electrically insulating fibrosis (structural repair tissue which is generated when the heart is damaged) infiltrates the space in between the electrically conducting heart muscle cells, see Fig. 4.4.

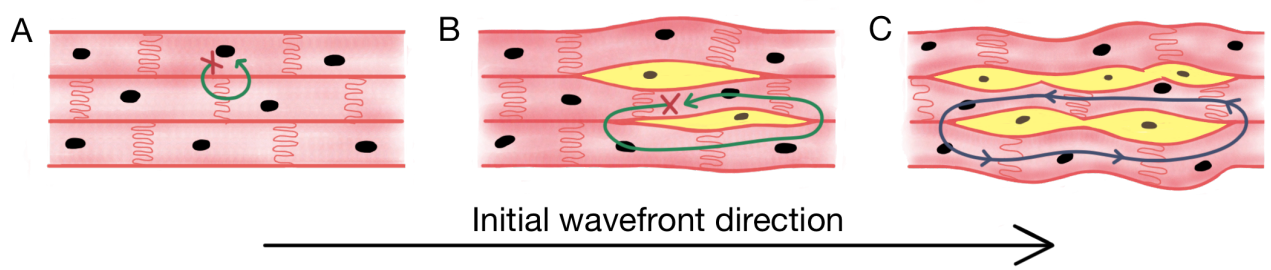

Figure 4.4: A schematic showing the formation of a micro-anatomical reentrant circuit from the accumulation of interstitial fibrosis. (A) In the healthy atria, heart muscle cells (pink cells), which are cylindrical in shape, are well coupled along and across the principle fibre directions. Therefore, no wavefront can back-propagate and form a circuit without colliding with its own unexcitable wavetail (green arrows blocked by red crosses). (B) Over time, the atria accumulate electrically insulating interstitial fibrosis (yellow cells) forming short segments of electrically isolated fibre. (C) The length of these segments grows as more interstitial fibrosis accumulates until the isolated segments are sufficiently long to harbour micro-reentry (blue arrows).

As insulating fibrosis builds up in the atria, isolated electrical channels form which, if sufficiently long (such that a wavefront does not collide with its own unexcitable wavetail), allow the smooth electrical wavefronts passing across the atria to back propagate and generate new, high frequency fibrillatory waves. The circuits which form via this process are often so small that, in

clinical practice, the circuit may appear as a point source, and may be one possible cause of the rapid ectopic triggers that are observed in the pulmonary vein sleeves. More detail regarding the electrical cycle of heart muscle cells and the formation of reentrant circuits is provided in section 5.8.

At its core, the hypothesis we pursue in the coming chapters is that the formation, frequency, and distribution of micro-anatomical reentrant circuits can be understood by framing the accumulation of fibrosis in the atria as a bond percolation problem. In such a framework, healthy tissue can be thought of as a fully connected lattice of cells, where each lattice site represents a group of heart muscle cells, and a bond between lattice sites represents strong electrical coupling between muscle cells in a local area. Then, if insulating fibrosis accumulates in the atria, this is equivalent to lowering the bond occupation probability on our lattice. Under such a construction, the formation of new micro-anatomical reentrant circuits is associated with the bond occupation probability gradually approaching the percolation threshold from above in our lattice or network. In chapter 5, we study the percolation approach on a 3d cubic lattice, whereas in chapter 6 we extend this to a more flexible spatial network approach.

The aim of this part is to explore the conceptual and practical utility of the percolation framework for AF. Summarising, the framework gives insights into how micro-anatomical reentrant circuits may distribute throughout the atrial wall based on the thickness of the muscle and the orientation of muscle fibres. Specifically, our results indicate that as AF becomes more severe, circuits are far more numerous and not very clustered, in contrast to mild cases where a small number of circuits are found in a confined area. In my opinion, this implies that the efficacy of ablation as a treatment for AF may have reached its practical limit.

Chapter 5

A percolation model of atrial fibrillation

5.1 Publication

Chapter 5 is based on the following publication:

M. Falkenberg *et al.* "Unified mechanism of local drivers in a percolation model of atrial fibrillation". In: *Physical Review E* 100.6 (2019), p. 062406. DOI: 10.1103/PhysRevE.100.062406

Contributions: The ideas in this chapter were developed by myself and Kim Christensen, based on earlier work in [145]. The text was written by myself with comments from Kim Christensen and Nicholas S Peters. Code used to generate the results in this chapter was written by Andrew J. Ford, Anthony C. Li and Robert Lawrence, supervised by myself, Kim Christensen and Alberto Ciacci. All figures were prepared by me and Andrew J. Ford.

5.2 Introduction

Atrial fibrillation (AF) is the most common cardiac arrhythmia [146], but our understanding of the underlying mechanisms is still poor [138, 139, 147–152]. There is growing evidence in

favour of local drivers as a sustaining mechanism of AF [144, 148, 152–160]. However, there is debate about the mechanistic origin of local drivers: Some studies have identified the drivers as rotors (reentrant drivers) [155, 156], whereas others have identified the drivers as focal points [154].

Recently, a study with "the potential to unify some of the previous discrepant observations" [161] has suggested that both focal and reentrant drivers may be explained by the presence of small micro-anatomical reentrant circuits [142, 144]. Using high-resolution simultaneous optical mapping of the endocardium (inner heart wall) and epicardium (outer heart wall) in explanted, diseased human hearts, the study shows that stable transmural reentrant circuits may project differently onto the endo- and epicardium. Projections typically appear as rotational activity on the endocardium and focal (breakthrough) points on the epicardium. Hence, the apparently incompatible 2d projections onto the atrial walls are consistent with a single underlying mechanism of localised reentrant circuits in the transmural region [144, 157, 158].

Leading clinicians have called for further work "to confirm and extend these observations"[161]. In this chapter, we develop a percolation based model of AF initiation in pursuit of this goal. The model we introduce is the 3d extension of the 2d cellular automata model of AF introduced by Christensen *et al.* [145].

Our aim is to understand the origin of the AF mechanism, how it changes as AF becomes more or less persistent, and how this affects the success rate of ablation. We take a physics approach to modelling, focusing on the essential features of the phenomenon as opposed to the fine details. This is much simpler than more biophysically realistic models and allows for concrete predictions based on large scale statistics not accessible in a laboratory setting.

The model generates activation wavefronts that propagate in a 3d medium mimicking the complex discrete fibre structure of the atria. AF emerges spontaneously via the formation of spatially stable but temporally intermittent 3d reentrant circuits when the fibre network decouples, e.g., through fibrosis or fatty infiltration [162]. The model predicts that these circuits should have a minimum length of 12.5 mm, comparable with the 15 mm quoted in Ref. [144]. Reentrant circuits are quasi-1d with isolated fibres that have a width of around 2 mm, making

their identification with clinically available mapping technologies difficult [163].

Our results demonstrate how 3d reentrant circuits emerge and how they are distributed in the heart wall when the coupling between cells is lowered towards the percolation threshold, closely matching key clinical observations in Ref. [144]. The observed activation patterns fundamentally change as a function of the depth from the imaged surface to the driving reentrant circuit. If this depth is large, we observe focal (breakthrough) activity. If this depth is small, we observe reentrant or rotational activity.

The model predicts that the drivers of paroxysmal AF are located near the endocardium whereas the drivers of persistent AF are uniformly distributed through the atrial wall. Consistent with clinical experience, paroxysmal AF is more easily treated by catheter ablation than persistent AF. We test this by numerically simulating focal ablation. We find that ablation lesions that do not penetrate the full depth into the tissue are significantly less likely to destroy a reentrant circuit in persistent AF than in paroxysmal AF. Additionally, we show that in persistent AF, even if ablation successfully destroys a local source, the existence of multiple co-existing circuits prevents the global termination of AF.

5.3 Model definition & properties

Electrical signals in the heart are mediated by discrete heart cells arranged in long, intertwined fibres, as shown in Fig. 4.4 in the introduction to this part. This motivates the construction of a 3d lattice model where each node acts as a single cell (or block of cells), and nodes are connected stochastically to their neighbours to mimic the branching structure of discrete heart muscle fibres, generalising a 2d model [145].

Our model simulates the initiation and maintenance of AF from local micro-anatomical reentrant circuits, which, due to their size, may appear as a focal trigger of AF. We do not investigate local drivers linked to cardiomyocyte automaticity, or the maintenance of fibrillation through the formation of rotors or scroll waves. Such phenomena are best studied using continuous reaction-diffusion models of cardiac electrophysiology [164].

Discrete models of cardiac tissue are well established historically and have recently re-gained popularity [145, 165–171]. In these models, the onset of re-entry has been associated with the approach from above of the bond occupation probability to the percolation threshold [168]. Some models have been extended to 3d to highlight the importance of the thickness of the atrial wall to the probability of re-entry [169]. Specifically when investigating the role of fibrosis in cardiac arrhythmia, which is the approach taken here, there have been suggestions that discrete models are preferable over continuous models [172]. Note, mechanistically it is important to distinguish between 3d vortices in a continuous substrate which exhibit functional re-entry and 3d reentrant circuits where electrically isolated fibres exhibit structural re-entry. The clinical results in Ref. [144] concern structural reentrant circuits and not the wider literature on rotors and scroll waves [164, 173].

We consider a simplified $L_x \times L_y \times L_z$ pipe topology of the atria with open boundaries in x and z, and periodic boundaries in y. Nodes are connected longitudinally to their neighbours in the x direction with spatial frequency $\nu_{\parallel}$ and transversely, in the y and z directions, with frequency $\nu_{\perp}$. Once the network is defined at the start of a simulation, it is fixed for the remainder of the simulation.

Each node takes one of three states: resting, where the node can be excited by an active neighbour, excited, which causes a heart muscle cell to contract, or refractory, where for τ time steps after excitation a node cannot be re-excited. We define the sino-atrial node as nodes at the boundary, $x = 0$, which excites every T time steps. Model parameters $L_x = L_y = 200$, $L_z = 25$, $T = 220$ and $\tau = 50$ are informed by clinically observed values, see section 5.8. The model results are robust against changes in these parameters. A small fraction of nodes, δ, are susceptible to conduction block. With probability ϵ, these nodes fail to activate when their neighbours excite.

When the sino-atrial node excites in normal conduction, wavefronts are initiated at $x = 0$ and propagate smoothly in the $+x$ direction and terminate at $x = L_x$. However, reentrant circuits can form when nodes are sufficiently decoupled, simulating the accumulation of diffuse interstitial fibrosis [168, 174], such that the shortest closed loop from a node back to itself is

partially isolated from the remaining tissue, and the path length, in units of propagation time, exceeds the refractory period τ. The formation of these reentrant circuits drive fibrillatory activity [6, 145], see Fig. 5.1.

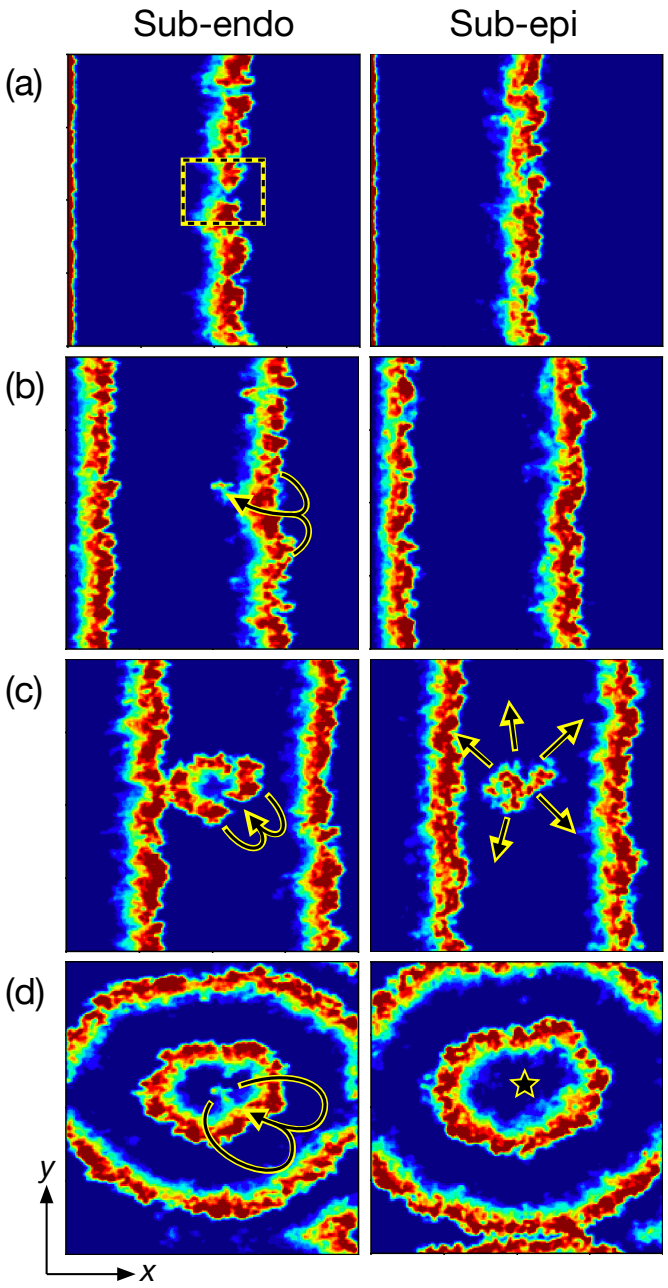

Figure 5.1: The emergence of AF in the percolation model with simultaneous endo- ($z = 0$) and epicardial ($z = 24$) imaging. Smoothing has been applied for clarity. Red: excited nodes, blue: resting (excitable) nodes, other: refractory (unexcitable) nodes. (a) Planar wavefronts propagating during sinus rhythm. Dotted box indicates gap in wavefront formed by conduction block. (b) Endo view (innermost layer): Arrow indicates excitation re-entering gap in wavefront. Epi view: Re-entry is not observed. (c) Emergence of fibrillatory activity. (d) Maintenance of fibrillatory activity. Epi view (outermost layer): Activity emerges on surface as point source (breakthrough activity) located at the star.

Ref. [144] finds that reentrant activity is typically observed on the endocardial surface, whereas focal points are often found on the epicardial surface. Additionally, there is much stronger longitudinal coupling on the endo- than the epicardial surface. To test whether fibre orientation can account for the distribution of activation patterns observed clinically, we consider a homogeneous and inhomogeneous model.

In the homogeneous model, all nodes are connected to their neighbours with the same frequencies, $\nu_{\parallel}$ and $\nu_{\perp}$. In the inhomogeneous model, the variation in fibre direction changes with depth. Here we fix the average decoupling of nodes, $\bar{\nu} = (2\nu_{\parallel} + 4\nu_{\perp})/6$ and vary linearly the average fibre angle, $\Delta\theta = \tan^{-1}(\nu_{\perp}/\nu_{\parallel})$, in each layer from $\Delta\theta_{\mathrm{ENDO}} = 24^{\circ}$ at the endocardium ($z = 0$) to $\Delta\theta_{\mathrm{EPI}} = 42^{\circ}$ at the epicardium ($z = 24$). These values are taken from clinical data [158].

The introduction of fibre inhomogeneity into the model results in reentrant circuits preferentially anchoring to the endocardium. An example of this can be seen in Fig. 5.2 which shows a cross-section of the model activation patterns from the endocardium to the epicardium. A reentrant circuit is shown anchored to the endocardium, with the activation wavefronts propagating from one end of the circuit to the epicardial surface and emerging as a focal breakthrough point.

To study the effect of local ablation in the model, we identify the first location at which a reentrant circuit forms. A new simulation with the identical structure is generated and the nodes at the identified location are destroyed, mimicking focal ablation, up to a prespecified depth, see section 5.9 for details. This process is repeated until the locations of the first 10 identified reentrant circuits are ablated. This allows us to study (a) the *global* success rate of ablation as a function of $\bar{\nu}$ by measuring the probability that ablation prevents AF emerging anywhere in the tissue, and (b) the *local* success rate of ablation as a function of the ablation depth z, by measuring the probability that an ablation, that does not penetrate the full depth into the tissue, destroys the reentrant circuit at a given location.

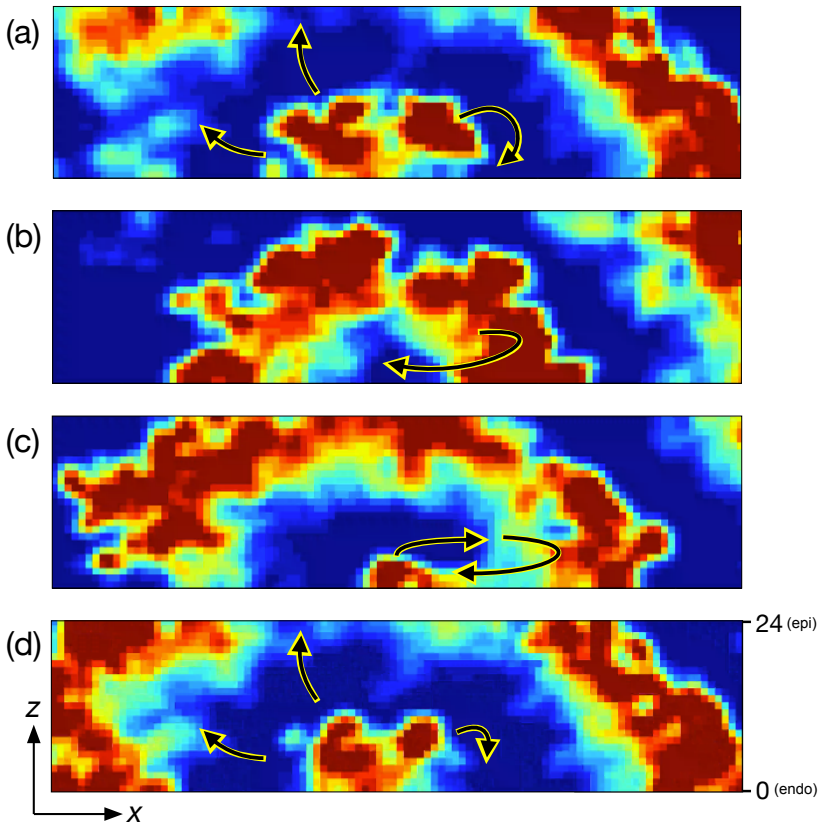

Figure 5.2: A cross-section of the inhomogeneous model during AF showing a close-up of the x-z plane (fixed y coordinate). The cross-section has been chosen to align with a reentrant circuit. Arrows indicate the movement of wavefronts. The top (bottom) surface of each panel corresponds to the epicardium (endocardium). A reentrant circuit can be seen near the endocardium with (a) activity moving around an isolated fibre, (b/c) re-entering the isolated fibre, and (d) emitting fibrillatory waves from the left end of the isolated fibre. Fibrillatory activity spreads from the endocardium to the epicardium through the complex fibre structure of the model. This results in a wide variety of possible breakthrough patterns on the epicardial surface. When viewed from the endocardium, reentrant activity will be clearly visible since the isolated fibre lies along the endocardial surface. When viewed from the epicardial surface, the wavefronts propagating from the reentrant circuit may emerge as a single breakthrough point, or as multiple breakthrough points simultaneously. Scale: 100×25 nodes.

5.4 Results

The phase spaces for the risk of entering AF for the homogeneous and inhomogeneous models are shown in Fig. 5.3. For large values of the coupling parameters (no fibrosis) the model exhibits sinus rhythm indefinitely. As the coupling reduces, a small number of reentrant circuits can form. Here we observe paroxysmal AF with intermittent episodes of irregular activity. As the coupling is even further reduced, the model enters persistent AF where, once AF has been initiated, the model will never return to sinus rhythm without external intervention. These

results are consistent with recent evidence showing that local drivers anchor at or near fibrotic lesions [156], with fibrosis increasing the number of reentrant regions, and the time spent in AF [175].

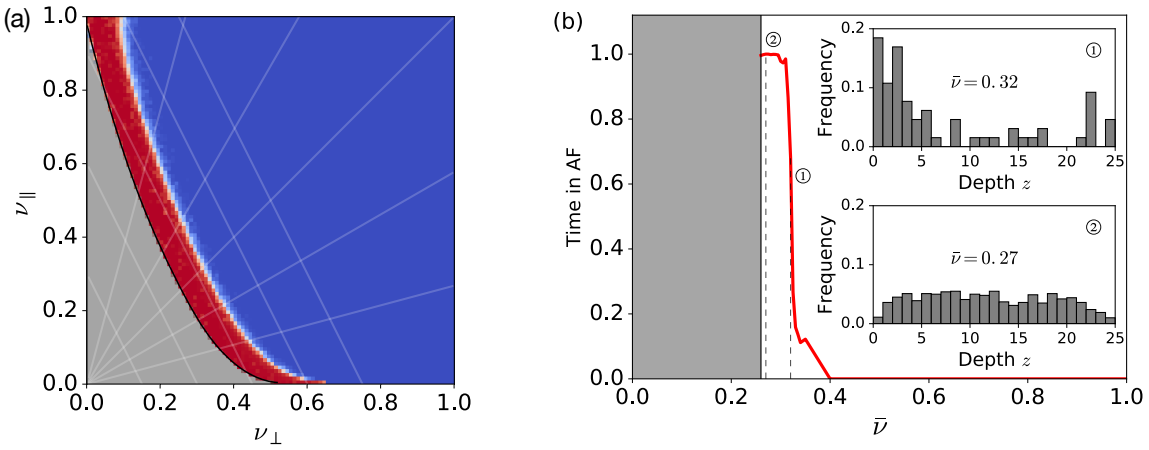

Figure 5.3: The average time spent in AF for the (a) homogeneous and (b) inhomogeneous models. Both models show a transition from sinus rhythm at large coupling (0% AF time) through paroxysmal AF (0% < AF time < 100%) to persistent AF (100% AF time). The grey parameter regions, indicate coupling parameters below the percolation threshold, and are thus irrelevant to clinical AF. (a) Homogeneous coupling parameters, $\nu_\parallel$ and $\nu_\perp$. Blue: sinus rhythm, red: persistent AF. The white transition region corresponds to paroxysmal AF. Guidelines of positive (negative) gradient indicate constant $\Delta\theta$ ($\bar{\nu}$). (b) Inhomogeneous coupling parameter, $\bar{\nu}$, risk curve (red graph). There is no AF for $\bar{\nu} > 0.4$. Decreasing $\bar{\nu}$ associated with decoupling nodes, there is a transition from sinus rhythm through paroxysmal AF to persistent AF. Inset: Histograms showing the distribution of reentrant circuits driving AF for paroxysmal AF (top) and persistent AF (bottom) as a function of depth z from the endocardium. Drivers cluster in the sub-endocardial region ($z = 0$) for paroxysmal AF but are uniformly distributed across the bulk for persistent AF.

In the homogeneous model, the endo- and epicardium are equivalent, and we find that in paroxysmal AF, the majority of drivers form equivalently on either surface, whereas for persistent AF, the majority of drivers form in the bulk and not on the surfaces, see Fig. 5.9 in section 5.10.

The inhomogeneous model breaks the symmetry between the endo- and epicardium. Here we find that in paroxysmal AF, drivers form preferentially near the endocardial surface, with very few drivers forming on the epicardial surface and almost none in the bulk of the atria, see Fig. 5.3(b), inset 1. However, in persistent AF, drivers are uniformly distributed throughout the atrial wall, see Fig. 5.3(b), inset 2. Hence, as AF becomes more persistent, the average position of drivers moves away from the endocardium and into the bulk of the atrial wall. The

variation in fibre orientation is the only asymmetry between the endo- and epicardium and is therefore responsible for the asymmetry in the reentrant circuit depths. Because the method by which reentrant circuits form is fully local [176], the presence of an isolated fibre at one point in the tissue is independent of the tissue elsewhere in the model. Therefore, the results are robust against different choices of fibre orientation and tissue thickness that may be associated with the left or right atria [177], and show that reentrant circuits will first emerge on the surface with the strongest longitudinal coupling.

Visualising the activation patterns in the inhomogeneous model on the surfaces at $z = 0$ (endocardium) and $z = 24$ (epicardium), we observe activation patterns consistent with the experimental results in Ref. [144]. In Fig. 5.4, the left (right) panel shows the activity when viewed from the endocardium (epicardium).

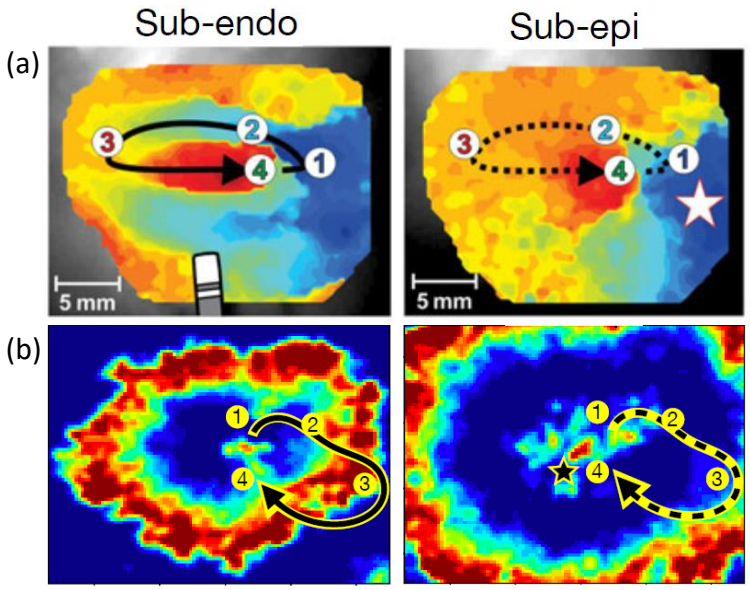

Figure 5.4: (a) An experimentally acquired activation map adapted with permission from Ref. [144] showing reentrant activity from sub-endo imaging and a focal source on the sub-epi side with four marked reference points. (b) The equivalent for our model. Smoothing has been applied to images for clarity. Red: excited nodes, blue: resting nodes, other: unexcitable refractory nodes. A pdf copy of the rights agreement allowing the reuse of panel (a) is appended to the end of this document, courtesy of Oxford University Press.

For the activation maps, electrical activity spreads across the tissue from point (1) through point (2), but is blocked from reaching the isolated fibre at point (4). The activity loops around the isolated fibre and re-enters the fibre at point (3). This back-propagating excitation passes

through point (4) before re-exciting point (1) which sustains the reentrant circuit. Viewing the same region from the epicardium does not show the same reentrant activity. Instead, the excitation emerges as a point source indicated by a star. The surface activity driven from the atrial bulk will typically appear as a breakthrough point, as opposed to reentrant activity. A cross-section of the activation patterns from endo- to epicardium is shown in Fig. 5.2.

Figure 5.5 shows the global and local success rate of ablation in the model. Globally, Fig. 5.5(a) indicates that the success of ablation exhibits phase transition-like behaviour.

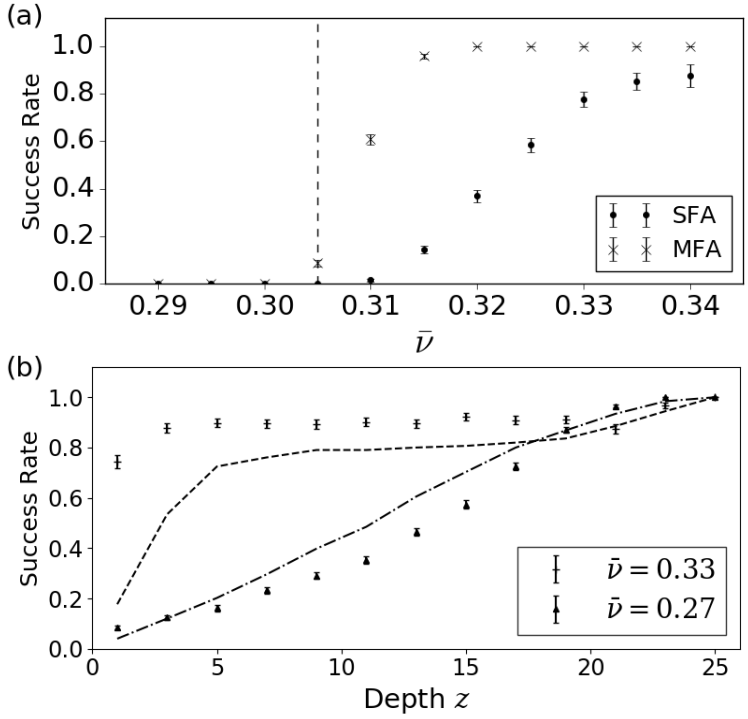

Figure 5.5: (a) The global success rate of full depth focal ablation if the first reentrant circuit location is destroyed (single focal ablation (SFA), circles), and if the first ten locations are destroyed (multiple focal ablations (MFA), crosses) as a function of the coupling parameter, $\bar{\nu}$. The dashed line indicates the coupling value below which all simulated tissues harbour reentrant circuits. (b) The local success rate of a single focal ablation of destroying reentrant circuits at the ablation site as a function of the ablation depth, z, for tissues with high coupling, $\bar{\nu} = 0.33$ (plus) and low coupling, $\bar{\nu} = 0.27$ (down triangle). The dashed (dot-dashed) lines correspond to the cumulative sum of the reentrant circuit depth probabilities for $\bar{\nu} = 0.33$ ($\bar{\nu} = 0.27$).

At low coupling, focal ablation consistently fails to terminate fibrillation, independent of the number of ablation lesions. Above a transition coupling value, the ablation success rate rapidly increases with increased coupling. The success rate is higher if multiple ablations are applied. These results indicate that focal ablation is increasingly ineffective as AF becomes more per-

sistent. However, despite the fact that ablation fails globally, Fig. 5.5(b) indicates that locally, ablation does successfully destroy the targeted reentrant circuits if the ablation is transmural.

For non-transmural ablation, ablation has a much higher local success rate at high coupling values – half depth ablation has a success rate of 90% (50%) at high (low) coupling. A visual comparison between transmural and non-transmural ablation is shown in Fig. 5.6.

The activation maps shown demonstrate that for reentrant circuits anchored to the endocardium, a shallow ablation is sufficient to destroy the reentrant circuit. However, for reentrant circuits anchored to the epicardium, non-transmural ablation may fail to destroy the reentrant circuit driving AF.

In Fig. 5.3(a), there are percolation-type transitions [179] associated with the black and white curves, respectively. The black curve outlines the bond percolation threshold for the fibre network itself: Below (above) the black curve, the fibre network does not (does) percolate from $x = 0$ to $x = L_x$. This transition is irrelevant for clinical AF but the other transition is highly relevant. The white curve outlines this transition where the finite surfaces associated with missing bonds (known as hulls in 3d or holes in 2d [180]) have a linear dimension larger than or equal to $\tau/2$ - these hulls are needed to form an isolated fibre. A fingerprint of the latter transition is also seen in Fig. 5.5(a). Below (above) the vertical dashed line at $\bar{\nu} \approx 0.305$, the global success rate of multiple focal ablation is zero (non-zero). In an infinite lattice, this corresponds to the number of finite hulls with linear dimension $> \tau/2$ being infinite (finite) slightly below (above) this threshold.

5.5 Discussion

These model observations are directly compatible with recent clinical findings that identified rotational activity from endocardial mapping in patients with paroxysmal AF [155, 176], and another clinical observation identifying focal activity in patients with persistent AF from epicardial mapping [154]. Moreover, our model offers a natural explanation why ablation is more successful for paroxysmal than persistent AF. Studies have shown that ablation lesions become

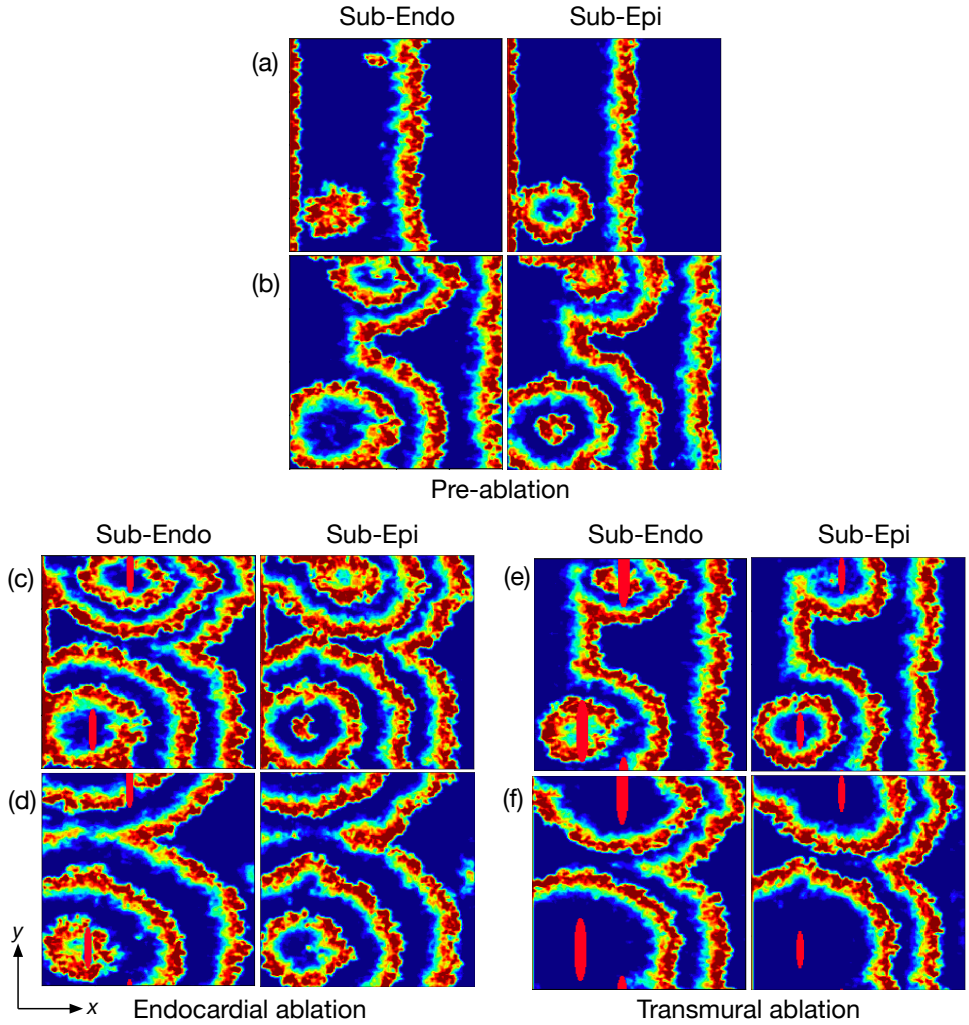

Figure 5.6: Simulated ablation in a tissue with multiple drivers demonstrating how the ability to terminate fibrillatory activity is closely related to driver depth. (a) AF emerges in the bottom left corner from a reentrant circuit near the epicardium, and (b) at the top of the tissue from a separate reentrant circuit near the endocardium. (c) Two shallow ablations (solid red in bottom left and top centre) are applied to the endocardial surface at the locations of the reentrant circuits penetrating part of the way into the tissue. (d) The circuit near the epicardium has not been destroyed by the ablation since the ablation has not penetrated far enough into the tissue. The circuit near the endocardium has been destroyed successfully. However, since an active circuit remains in the tissue, AF has not been terminated. (e) & (f) show the same ablation process as demonstrated in (c) & (d), however, the ablation is now transmural, penetrating the full distance through the tissue. This is indicated by the solid red regions on the sub-epi view in (e) & (f). The ablations have successfully destroyed both reentrant circuits. Hence, fibrillatory waves dissipate as shown in (f) and sinus rhythm will be restored. Note, transmural ablation is not always possible in clinical practice. Ablation lesions typically struggle to penetrate more than 2 mm [178] into the atrial wall which can be up to 7 mm thick in places [144]. Scale: 200×200 nodes (full tissue).

smaller as they penetrate further into the tissue from the endocardial surface, and that ablation struggles to penetrate more than 2 mm into the atrial wall [178] which can be up to 7 mm

thick [144]. Hence, because the average driver position moves deeper into the atrial wall as AF becomes more persistent, ablation lesions need to be more accurately positioned and they must penetrate further into the tissue, restricting ablation efficacy. The model suggests that if the atria are sufficiently fibrotic, reentrant circuits are so numerous that focal ablation cannot terminate AF, even with multiple ablation lesions, see Fig. 5.5(a).

The model may explain the prevalence of paroxysmal AF in the pulmonary veins (PVs), where sleeves of cardiac tissue extend into the PVs [142, 143]. The sleeves get thinner further into the PVs and there are significant changes in the fibre orientation from the endo- to epicardium [181–183]. These physiological changes reduce cell-to-cell coupling which the model has shown is the key requirement for forming the reentrant circuits that drive AF.

5.6 Limitations & next steps

Our model is much simpler than alternative biophysically realistic models [164, 173]. We do not account for rate dependent effects in the action potential propagation, although we expect this to play a minimal role in the development and maintenance of structural reentrant circuits. This has the benefit that we are able to perform unparalleled statistical analyses to fully explore the differences between paroxysmal and persistent AF, but at the cost of electrophysiological realism. As a result, our model focuses purely on the initiation and maintenance of AF from micro-anatomical reentrant circuits and not on AF initiated or maintained by other methods such as cardiomyocyte automaticity or the formation of rotors and scroll waves. These mechanisms are best studies using alternative models.

In our model, we consider general structural properties such as the difference in fibre orientation between the endo- and epicardium, but we do not explicitly account for specific anatomical fibre bundles in the real heart or differences between the left and right atria.

To address these limitations and verify the results presented here, we propose a potential experimental test and a number of suggested model developments. Firstly, adapting the model to a realistic atrial topology with an integrated fibre structure will allow us to refine the predic-

tions made here and make concrete statements about different regions of the atria and specific anatomical structures. This includes investigating which atrial regions are especially prone to the formation of reentrant circuits and testing whether these correlate with known clinical data. Chapter 6 aims to address these limitations and answer some of these questions.

5.7 Chapter summary

- We introduce a 3d lattice-based model of atrial fibrillation (AF), where the propagation of electrical wavefronts is implemented using cellular automata style rules.

- The model, which considers the percolation-based decoupling of a 3d lattice to simulate the accumulation of electrically insulating fibrosis, is designed to study the formation of micro-anatomical reentrant circuits. These circuits have been implicated as a potential mechanism for initiating and maintaining AF.

- Our aim studying the percolation model is to understand how the distribution of micro-anatomical reentrant circuits is affected by the degree of coupling in the lattice. We find a rapid increase in the risk of micro-reentry as the lattice decouples, spanning the observed spectrum of AF risk from sinus rhythm, through paroxysmal AF, to persistent AF.

- In paroxysmal AF, circuits are rare and cluster near the atrial surfaces. In persistent AF, circuits are numerous and distribute evenly across the atrial bulk.

- In one version of the percolation model, we introduce variable fibre orientation between the inside and outside of the 3d lattice, representing the endo- and epicardial surfaces of the atria. We find that stronger longitudinal coupling on the endocardium results in circuits preferentially clustering near the endocardium, rather than at the epicardium.

- AF is often treated by catheter ablation, destroying the atrial tissue in regions believed to initiate and maintain AF. We show that the probability of carrying out a successful ablation reduces rapidly as AF becomes more persistent.

5.8 Technical details

5.8.1 Model definition

Heart muscle structure

The area of the atrial muscle is $L_x^{'} \times L_y^{'} \approx 20\text{cm}^2$ [184, 185] while the thickness of the atrial muscle wall ranges from 1 mm to 7 mm [144]. A typical average is about $L_z^{'} = 2.5\,\text{mm}$ [186]. The atrial muscle tissue comprises tubular shaped heart muscle cells (myocytes) of length $\Delta x^{'} \approx 100\mu\text{m}$ and diameter $\Delta y^{'} = \Delta z^{'} \approx 20\mu\text{m}$ [187, 188]. Cells are connected longitudinally and transversally through gap-junctions which mediate the flow of ions between cells. The density of gap-junctions is high at the ends of a cardiomyocyte, and lower along the cardiomyocyte. Hence a branching network of anisotropic cells provides an appropriate approximation of the atrial tissue [145]. A detailed description of cardiac structure is given in [137].

Heart muscle physiology

The membrane potential of heart muscle cells shows that they can be in one of three states: excited, where the cell depolarises, resting, where the cell can be depolarised by a neighbouring active cell, and refractory, where for some time τ' after excitation the cell cannot be re-excited. The shortest time scale is associated with a cell depolarising, $\Delta t' \approx 0.6\,\text{ms}$, whereas the refractory period lasts $\tau' \approx 150\,\text{ms}$ under rapid pacing. When a resting cell is induced to excite by a neighbouring excited cell, it contracts. The propagation speed of the activation wavefront is $\Delta x'/\Delta t' \approx 0.2\,\text{ms}^{-1}$ and the heart's natural pacemaker (sino-atrial node) ensures the periodic propagation of a coordinated activation wavefront, causing atrial muscle contractions. The rate of excitations are controlled by the pacemaker cells that regularly self-excite with period $T' = 300\text{-}1000\,\text{ms}$. For example $T' = 660\,\text{ms}$ corresponds to approximately 90 beats per minute.

Dimensionless parameters of the model

We can translate real tissue values into dimensionless model parameters. We apply a similar conversion procedure to that presented in Ref. [145]. The dimension of the simplified 3d topology are $L_x = L'_x/\Delta x' = 1000$, $L_y = L'_y/\Delta x' = 1000$ and $L_z = L'_z/\Delta z' = 125$. We apply a coarse grain transformation $\Delta x' \to b\Delta x', \Delta y' \to b\Delta y', \Delta z' \to b\Delta z'$ such that we group b^3 cells within one node of the model. The propagation speed has to remain constant $\Delta x'/\Delta t' \approx 0.2\,\mathrm{ms}^{-1}$ so we must also coarse grain time $\Delta t' \to b\Delta t'$ [145]. In the current work, we set $b = 5$. For the 3d pipe structure, we have $L_x = L'_x/(b\Delta x') = 200, L_y = L'_y/(b\Delta y') = 200$ and $L_z = L'_z/(b\Delta z') = 25$. The renormalised excitation time $b\Delta t' = 3\,\mathrm{ms}$ yields a dimensionless refractory period $\tau = \tau'/(b\Delta t') = 50$ and pacemaker activation period $T = T'/(b\Delta t') = 220$, assuming $T' = 660\,\mathrm{ms}$.

Modelling the real atrial tissue

In the human heart, electrical signals are initiated at the sino-atrial node and traverse through the mesh of cardiomyocytes to the atrio-ventricular node where the signal is transmitted from the atria to the ventricles. In the atria, electrical signals pass along cardiac muscle fibres formed by long strands of cardiomyocytes (individual muscle cells). Electrical conduction between these cells is mediated by the transfer of ions through gap-junctions. The density of gap junctions is high at the longitudinal ends of a cell, and lower in the transverse direction. This results in a natural longitudinal-transverse electrical conduction anisotropy in cardiac muscle fibres. As fibrosis accumulates the density of gap-junctions is reduced, particularly in the transverse direction, exacerbating conduction anisotropy [189].

The structural properties of the atria are mimicked by a 3d pipe-like topology, $L_x \times L_y \times L_z$, where x is the longitudinal direction and y and z the transverse directions. Given that the mechanism of reentrant circuit formation in the model is fully local, the precise topology chosen for the model has a minimal effect on the results.

The pacemakers cells are the nodes at the boundary $x = 0$. We apply open boundaries in x

and z and periodic boundaries in y. Each of the $L_z = 25$ layers contain $L_x \times L_y = 200 \times 200$ nodes. The nodes are connected longitudinally with frequency $\nu_{\parallel}$ and transversally (within the layer and across the layers) with frequency $\nu_{\perp}$. The network of nodes mimics the grid-like arrangement of cardiomyocytes in the heart which are well connected longitudinally, through the high density of gap-junctions, but have reduced transverse coupling, due to naturally lower transverse gap-junction densities and due to the accumulation of interstitial fibrosis.

To mimic the membrane potential of a cell, nodes can be in one of three distinct states: resting, excited or refractory, see Fig. 5.7. At a given time t, an excited node induces resting nodes to which it is coupled to excite at time $t+1$. The excited node at time t becomes refractory at time $t+1$ and it will not be able to respond to any further stimulus during the next τ time steps. The refractory node will revert to a resting node after τ time steps.

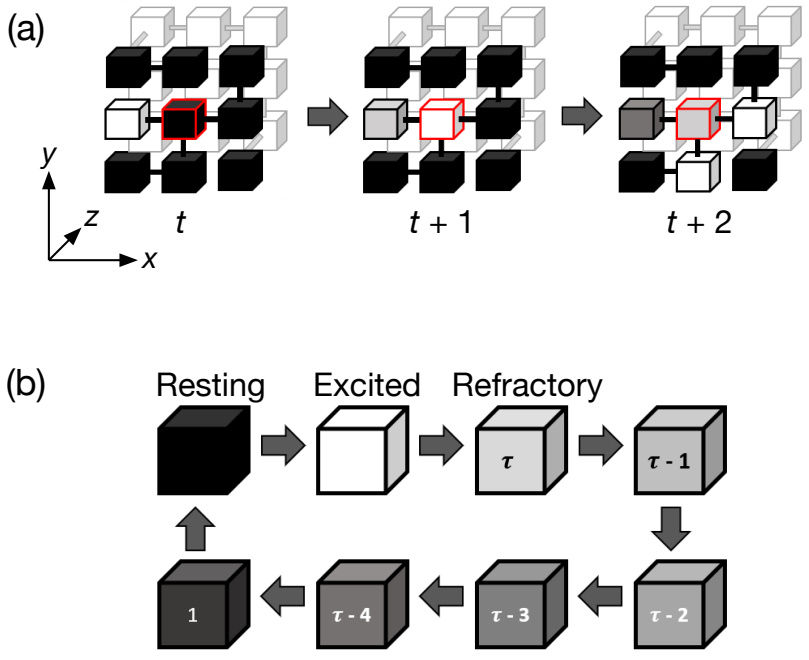

Figure 5.7: (a) The excitation wave propagates across neighbouring nodes. Nodes are connected longitudinally with frequency $\nu_{\parallel}$ and transversally (along y and z directions) with frequency $\nu_{\perp}$. At time t, the red bordered node is resting. At time $t+1$, the node activates in response to the stimulus coming from its neighbour to the left. At time $t+2$, the red bordered node becomes refractory and the excitation wave activates the two nodes in resting state to which it is coupled. The node to the left of the red bordered node cannot be excited as it is in a refractory state. (b) Temporal dynamics of a node. Each node sequentially passes across three distinct states: resting (black), excited (white) and refractory (grey scale).

A small fraction δ of the nodes can be susceptible to conduction block. These nodes are selected

at random and their position in the network is fixed for a given realisation. These particular nodes will, with a small probability ϵ, fail to activate when their neighbours excite. In a well connected network, such a failure will not disrupt the planar propagation of the excitation wave and sinus rhythm prevails. However, if the network of heart muscle cells becomes too decoupled, for example due to fibrosis, then, when a failure occurs, the planar propagation of the excitation wave may be disrupted.

The break of the regular wave front may prompt the spontaneous emergence of circuits formed by wavelets of excitation leaking back through the refractory wake of the wave front [145]. This occurs if nodes in a certain region of the atrial muscle form a reentrant circuit, see Fig. 5.8. Reentrant circuits can emerge spontaneously in the atria due to structural changes during the lifetime of a patient as specific pathologies (e.g., fibrosis) might decrease the coupling between cells.

5.9 Simulation details

5.9.1 Ablation methods

To simulate the effect of ablation in the model we generate 500 independent tissues for each coupling value, $\bar{\nu}$. The random seeds used to generate these tissues are saved so that the tissues can be regenerated at a later stage. An additional seed is generated for each tissue which specifies the order in which nodes susceptible to conduction block fail to excite when prompted by an excited neighbour.

For the global success rate of ablation, a simulation is started in each tissue up to a maximum time of 10000 time steps. For simulations of length 10000 ($\sim$ 45 beats), with $\epsilon = 0.05$ and $T = 220$, if a structure exists capable of generating a reentrant circuit is present, the probability it activates during the simulation is $> 90\%$. If a reentrant circuit is detected during the evolution of the simulation, the simulation is stopped and the location of the circuit is recorded. Otherwise, we mark this structure as containing no reentrant circuits and move on to the next.

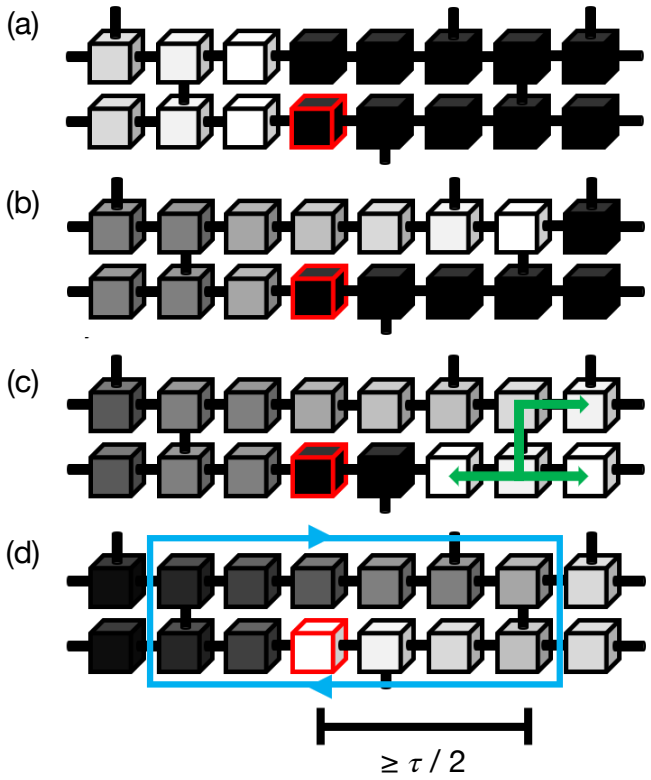

Figure 5.8: Progression of the excitation wave through a simple reentrant circuit situated on the surface of the 3d topology used to mimics the heart muscle fibre network. The red bordered node is susceptible to conduction block. (a)-(b) The planar excitation wave reaches the node that is susceptible to conduction block. The red bordered node does not excite and blocks the propagation in the lower fibre while the planar propagation of the wavefront proceeds in the upper fibre. (b)-(c) The activation wavefront reaches the penultimate node in the upper fibre that connects the two adjacent fibres. The excitation advances in multiple directions, indicated by the green arrows. In particular, the wavefront re-enters the fibre where the original conduction was blocked. (c)-(d) If the path length indicated in the lower fibre is longer than $\tau/2$, the node placed to the left of the red bordered node will have reached its resting state when the backward propagation arrives. Therefore, the region indicated in light blue will act as a reentrant circuit until the red bordered node blocks the excitation again. The likelihood of creating a reentrant circuit is determined by the coupling frequencies $\nu_{\parallel}$ and $\nu_{\perp}$ and the fraction δ of nodes susceptible to conduction block.

If a reentrant circuit is detected, using the saved seeds, the exact same tissue is regenerated and we reset the order in which nodes susceptible to conduction block fail to fire. Before the simulation is restarted, a cuboid of side length 20 nodes and depth 25 nodes, centred on the x and y co-ordinate identified as the circuit location, is destroyed, mimicking a focal ablation. For a tissue of dimensions $200 \times 200 \times 25$ nodes, this ablation lesion corresponds to 1% of the tissue. Destroyed nodes are permanently disabled and cannot be excited by neighbouring excited nodes. After the ablation lesion is applied, we restart the simulation for another 10000 time steps and observe whether another reentrant circuit emerges at a different location in the

tissue. If yes, this region can be ablated and the process can be repeated an arbitrary number of times – at this stage the focal ablation has been unsuccessful at terminating AF globally. If no, the ablation is considered successful and we move onto the next tissue.

The global ablation success curves shown in Fig. 5.5 correspond to the percentage success that ablation prevents the emergence of AF anywhere in the tissue after a single ablation has been applied, and after up to ten individual lesions have been applied.

For the local success rate of ablation, see Fig. 5.5(b), we also generate 500 random tissues for each coupling value and save their seeds. Likewise, we save the seeds giving the sequence of failures for the nodes susceptible to conduction block. For each tissue, we run a simulation for up to 10000 time steps. If a reentrant circuit is not identified we move onto the next tissue. If a circuit is identified, a cuboid of side length 20 nodes and depth z nodes, centred on the x and y co-ordinate identified as the circuit location, is destroyed. Additionally, all nodes susceptible to conduction block with coordinates *outside* the x and y extent of the ablation lesion are prevented from exhibiting unidirectional conduction block. Hence, when the simulation is restarted, a reentrant circuit can only form in the region directly below the ablation lesion, that is, a circuit can only form if it is at a co-ordinate corresponding to the x and y range ablated, but with depth greater than that of the ablation lesion. The simulation is now restarted for 10000 time steps and we observe whether or not a reentrant circuit is detected. If yes, the local ablation has failed. If no, the local ablation has been successful. This process is repeated for a range of ablation depths z from a single node layer to the full tissue depth $z = 25$, using identical tissues such that different depth results can be compared like for like.

There are three reasons local ablation might fail: (1) the ablation lesion is not deep enough to destroy a reentrant circuit, (2) the targeted reentrant circuit is destroyed but a different independent circuit is present below the ablation lesion, (3) the ablation lesion itself decouples nodes such that a circuit which was not present previously can now form by anchoring to the the ablation lesion itself.

5.10 Further model results

5.10.1 Depth of reentrant circuits

We can measure the depth from the endocardium (z coordinate) at which reentrant circuits form in the homogeneous and inhomogeneous model. Figure 5.9 (Fig. 5.10) shows the distribution of reentrant circuit depths for the homogeneous (inhomogeneous) models at various points on their respective phase spaces.

In the homogeneous model, the fibre structure is statistically equivalent across the whole tissue. Hence, there is no distinction between the endocardial surface at $z = 0$ and the epicardial surface at $z = 24$. Figure 5.9 shows that in paroxysmal AF, reentrant circuits form preferentially near the atrial surfaces. As nodes continue to decouple and the model approaches persistent AF, reentrant circuits form uniformly in the bulk. This dependence of reentrant circuit depth on AF persistence is found to exist for models with strong longitudinal and weak transverse coupling (Fig. 5.9 left column), equal longitudinal and transverse coupling (Fig. 5.9 central column), and weak longitudinal and transverse coupling (Fig. 5.9 right column). However, reentrant circuits cluster around the atrial surfaces more closely for strong longitudinal and weak transverse coupling, than for weak longitudinal and strong transverse coupling.

As a result, when the fibre orientation is varied in the inhomogeneous model, reentrant circuits will form preferentially in the regions with the strongest longitudinal coupling, see Fig. 5.10. This breaks the symmetry in reentrant circuit distribution between the endocardium and the epicardium in the inhomogeneous model. In the inhomogeneous model for paroxysmal AF, reentrant circuits are found to cluster preferentially at the atrial surfaces, as opposed to uniformly in the bulk, and at the endocardium, as opposed to the epicardium. As AF progresses, the asymmetry between epicardium and endocardium weakens and reentrant circuits move deeper into the bulk of the atrial wall. Hence, the average distance of reentrant circuits from the endocardial surface increases as AF becomes more persistent. As reentrant circuits move deeper into the atrial wall, they become harder to ablate.

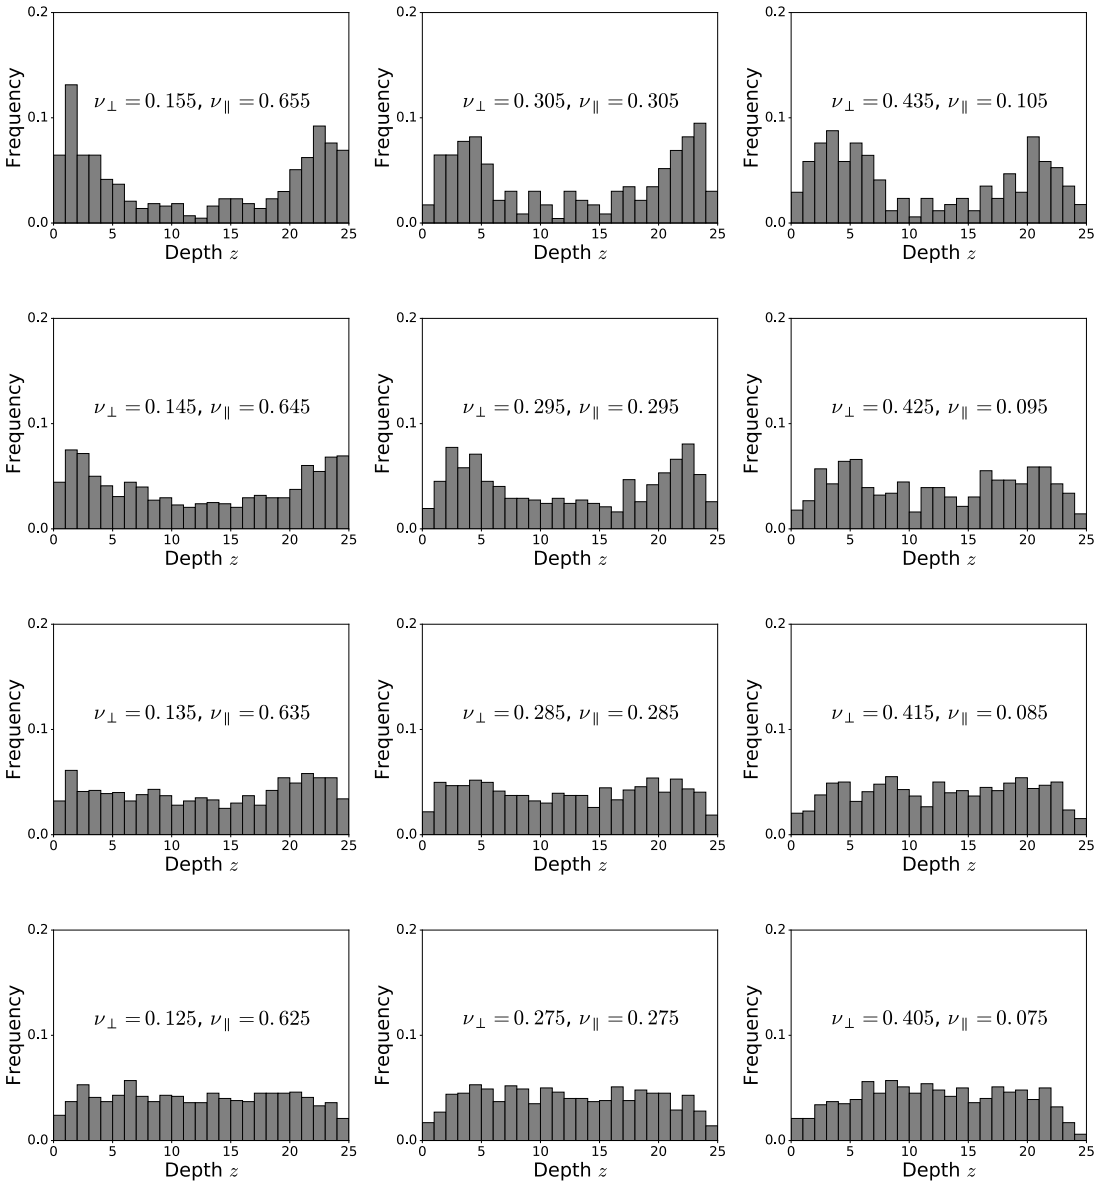
$\nu_\perp = 0.155$, $\nu_\parallel = 0.655$
$\nu_\perp = 0.305$, $\nu_\parallel = 0.305$
$\nu_\perp = 0.435$, $\nu_\parallel = 0.105$
$\nu_\perp = 0.145$, $\nu_\parallel = 0.645$
$\nu_\perp = 0.295$, $\nu_\parallel = 0.295$
$\nu_\perp = 0.425$, $\nu_\parallel = 0.095$
$\nu_\perp = 0.135$, $\nu_\parallel = 0.635$
$\nu_\perp = 0.285$, $\nu_\parallel = 0.285$
$\nu_\perp = 0.415$, $\nu_\parallel = 0.085$
$\nu_\perp = 0.125$, $\nu_\parallel = 0.625$
$\nu_\perp = 0.275$, $\nu_\parallel = 0.275$
$\nu_\perp = 0.405$, $\nu_\parallel = 0.075$
Frequency
Depth z

Figure 5.9: The distribution of reentrant circuit depth z from the endocardium in the homogeneous model. Fibre angle $\Delta\theta$ is constant across the tissue. AF progresses from paroxysmal AF (top row) to persistent AF (bottom row). The left column is for tissues with strong longitudinal coupling and weak transverse coupling. The central column is for tissues with equal longitudinal and transverse coupling. The right column is for tissues with weak longitudinal coupling and strong transverse coupling. Depth $z = 0$ ($z = 24$) corresponds to the endocardium (epicardium). In paroxysmal AF, reentrant circuits are predominantly found near the atrial surfaces. When visualised from the atrial surface, these circuits will typically appear as reentrant or rotational activity. In persistent AF, reentrant circuits are uniformly distributed in the bulk of the atrial wall, away from the epicardium and endocardium. These circuits are more likely to result in breakthrough point (focal) activation patterns when visualised from the atrial surfaces. In the homogeneous model, the endocardium and epicardium are equivalent due to the symmetry.

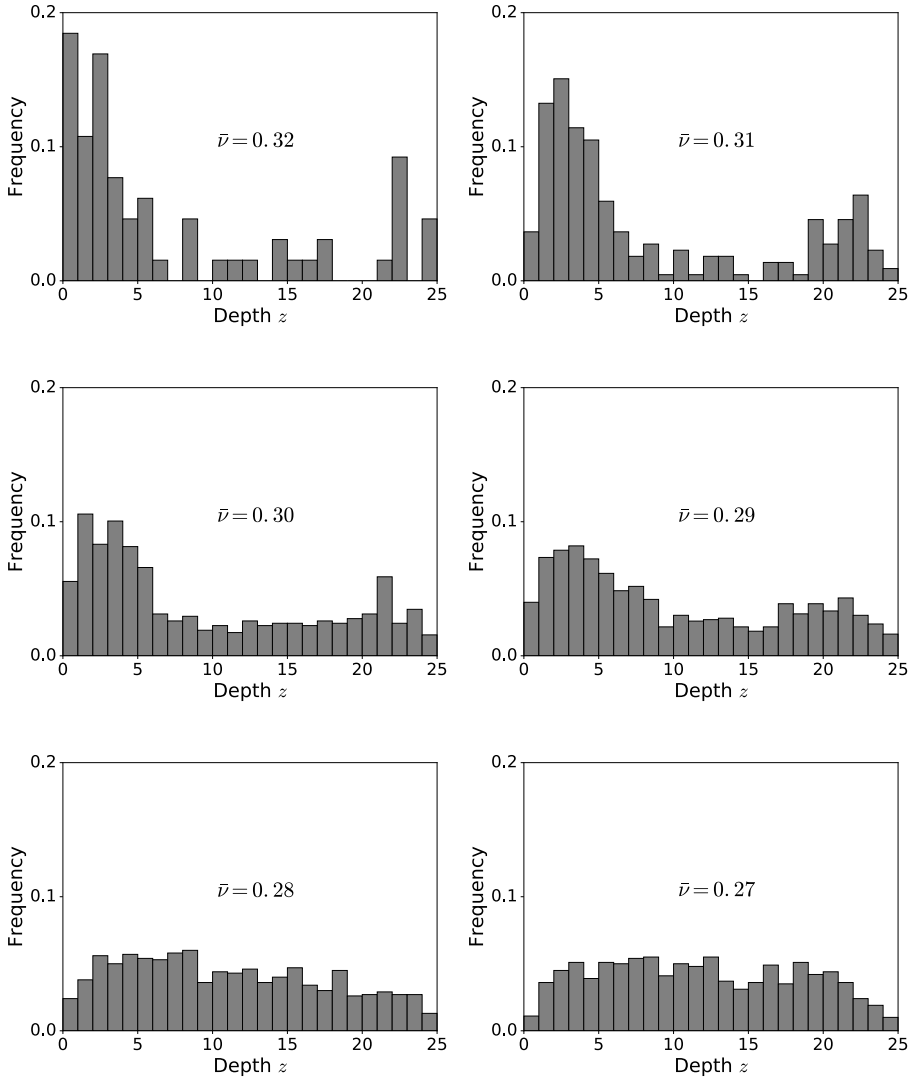

Figure 5.10: The distribution of reentrant circuit depth z from the endocardium in the inhomogeneous model. AF progresses from paroxysmal AF ($\bar{\nu} = 0.32$) to persistent AF ($\bar{\nu} = 0.27$). Fibre angle varies linearly across the tissue from $\Delta\theta_{\text{ENDO}} = 24°$ on the endocardium to $\Delta\theta_{\text{EPI}} = 42°$ on the epicardium. Depth $z = 0$ ($z = 24$) corresponds to the endocardium (epicardium). In the inhomogeneous model, the symmetry between the epicardium and endocardium is broken with the stronger longitudinal coupling on the endocardium than the epicardium. For paroxysmal AF, reentrant circuits are predominantly found near the atrial surfaces, and significantly more circuits cluster near the endocardium than the epicardium. This is because fibres have much stronger longitudinal coupling on the endocardial surface than on the epicardial surface. When visualised from the endocardium, these circuits will typically appear as reentrant or rotational activity. In persistent AF, reentrant circuits are uniformly distributed in the bulk of the atrial wall, away from the epicardium and endocardium. These circuits are more likely to result in breakthrough point (focal) activation patterns when visualised from the endocardium. As AF has progressed from paroxysmal to persistent, the average depth of circuits from the endocardium at $z = 0$ has significantly increased – this hinders the effective use of catheter ablation for those circuits deep in the atrial bulk.

Chapter 6

Identifying micro-anatomical reentry using spatial networks

6.1 Publications

Chapter 6 is based on the following two publications:

M. Falkenberg *et al.* "Identifying Potential Re-Entrant Circuit Locations From Atrial Fibre Maps". In: *2019 Computing in Cardiology (CinC)*. IEEE. (2019), pp. 1–4. DOI: 10.22489/CinC.2019.102

M. Falkenberg *et al.* "Identifying locations susceptible to micro-anatomical reentry using a spatial network representation of atrial fibre maps". In: *bioRxiv* (2021). DOI: 10.1101/2021.09.13.460069

Contributions: The ideas in this chapter were developed by myself and Kim Christensen. The text was written by myself and James A Coleman, with comments from Kim Christensen, Nicholas S Peters, Fu Siong Ng, Arunashis Sau and Jichao Zhao. Jichao Zhao and Belvin Thomas provided technical expertise and key datasets. The code used to generate our results was written by James A Coleman and Sam Dobson, supervised by myself, Alberto Ciacci and Kim Christensen. A previous proof of concept implementation of the current work was

developed by myself, Kim Christensen, Louie Terrill and David J Hickey. Figures were created by myself and James A Coleman, with the exception of the schematics in Figs. 6.1 and 6.2 which were drawn by Becky Joakim.

6.2 Introduction

In this chapter, we introduce a novel method to assess the feasibility of patient-specific predictions for the distribution of micro-anatomical reentry across the atria. In particular, the method is inspired by the idea that it may be possible to predict the emergence of micro-reentry before the atria have accumulated sufficient interstitial fibrosis, in contrast to most other patient specific approaches which take a static view of the AF substrate [190, 191].

Starting from image-based models of the atria, we combine data regarding the atrial geometry and the underlying myocardial fibre structure to form a spatial network. By progressively removing connections in the atrial structure, we assess where in the network a path exists which is sufficiently long to harbour micro-reentry. This technique is related to the study of spatial percolation in network science [192], and is the natural extension of the lattice-based approach in chapter 5. Percolation approaches have a strong history in computational electrophysiology having been used to explain the fractionation of electrograms [168] and the emergence of ectopic pacemakers [169].

In the current work, our primary goal is to understand the utility of the method and demonstrate its future potential. With this in mind, we ask (1) whether our method for constructing a spatial network preserves the underlying properties of the fibre orientation in each dataset, (2) how the predictions from our method depend on fibre structure and atrial geometry, and (3) whether the origin of the fibre data is important, i.e., whether both anatomically derived and synthetically generated fibre structures can be used.

In section 6.3 we detail our aims and describe the spatial network concept. In section 6.4 we introduce the image-based atrial datasets used in our study, and outline our spatial network generation and analysis methods. In section 6.5 we present our results, before discussing our

findings in section 6.6. Additional results and technical details are provided from section 6.9 onwards.

6.3 Aims

It has recently been shown "that human AF may be driven by micro-anatomic reentrant AF drivers anchored to fibrotically insulated tracks within the complex atrial wall" [193]. To ensure that a given fibre tract can sustain micro-reentry, the reentrant pathway through a fibrotically insulated fibre tract must be at least one refractory wavelength long.

Our approach is illustrated schematically in Fig. 6.1. Denoting the minimum refractory wavelength during fibrillation as $\tilde{\tau}$, our aim is to assess where in the atrial structure fibrotically insulated reentrant pathways exist of length ℓ, such that $\ell \geq \tilde{\tau}$. In a healthy heart with minimal interstitial fibrosis, atrial myofibres are well coupled in both the longitudinal and transverse directions. Hence, the longest possible reentrant path is significantly shorter than one refractory wavelength, $\ell \ll \tilde{\tau}$.

As the density of electrically insulating interstitial fibrosis increases, the longest reentrant pathway increases until a threshold density is reached at which $\ell \geq \tilde{\tau}$ and the fibrotically insulated fibre tract can sustain continuous reentry during AF. Importantly however, this threshold density may vary across different regions of the atria; some regions may be susceptible to micro-reentry at a low fibrosis density, whereas other regions may require a significantly higher fibrosis density before micro-reentry can be induced. If the density of fibrosis is too high in a region, micro-reentry may be prevented due to the absence of any closed loops (compact fibrosis). Hence, different atrial regions may have unique characteristic fibrosis density ranges in which micro-reentry is inducable.

Our hypothesis is that the process of electrically insulating fibre tracts with interstitial fibrosis can be modelled using spatial networks. A spatial network is a graph consisting of nodes representing entities in space, and edges representing spatial connections between those entities [194]. In our case, a node represents one or more atrial cardiomyocytes in a specific atrial region,

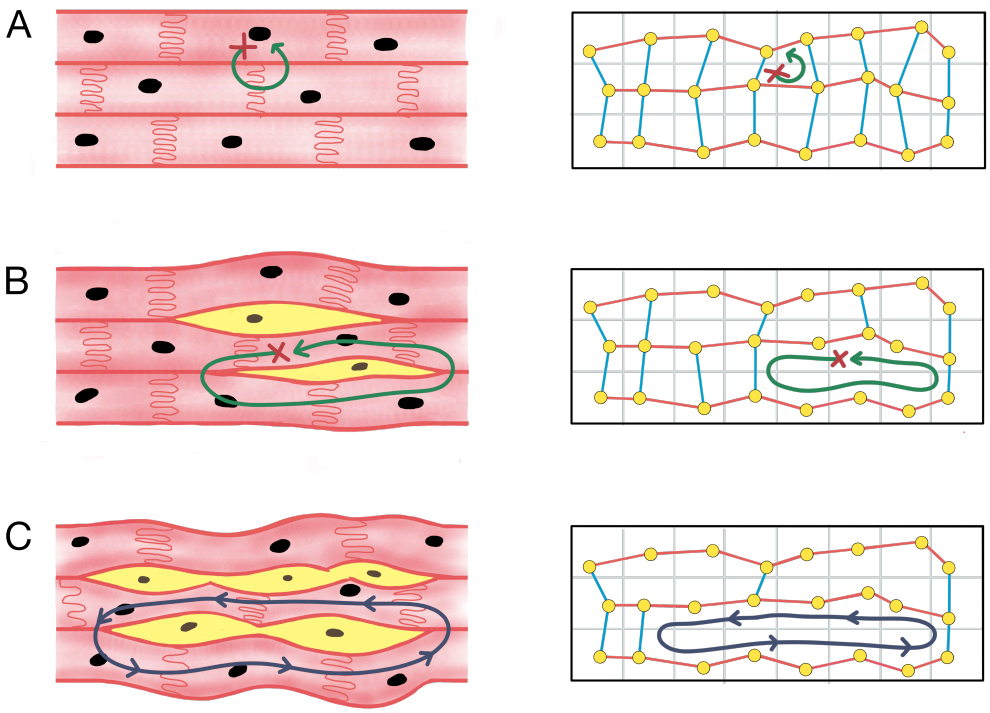

Figure 6.1: A schematic illustrating our approach. (A) In the healthy atrial myocardium, cardiomyocytes (pink cells) are well coupled along and across the principle fibre directions, such that the longest reentrant loop is too short to sustain micro-anatomical reentry (green arrows blocked by red crosses). (B) Over time, the atria accumulate interstitial fibrosis (yellow cells) forming short segments of electrically isolated fibre. (C) The length of these segments grows as more interstitial fibrosis accumulates until the isolated segments are sufficiently long to harbour micro-reentry (blue arrows). We hypothesise that the accumulation of interstitial fibrosis can be modelled as a spatial network. Right column: Nodes (yellow points) representing a group of cardiomyocytes are connected to their neighbours along fibres (red links) and across fibres (blue links). If the density of interstitial fibrosis is low, there are many transverse (blue) links. We model the increase in the density of interstitial fibrosis as the progressive removal of transverse links, and equate the micro-anatomical reentrant substrate to regions of the spatial network where a loop exists that is longer than one refractory wavelength. Figure drawn by Becky Joakim.

and an edge between two nodes indicates whether those groups of myocytes are electrically coupled, that is, whether an activation wavefront can pass directly from the group of myocytes represented by one node, to the group of myocytes at the other node. Naturally, the edges in the spatial network are short such that nodes can only be coupled to other nodes in their immediate vicinity.

Starting from a network in which the local density of edges is high, our assertion is that the accumulation of insulating interstitial fibrosis in the atria is structurally equivalent to the

progressive removal of edges in our spatial network. If the density of fibrosis is low, then the density of edges in the spatial network is high, and similarly, if the density of fibrosis is high, the density of edges in the spatial network is low. Hence, by identifying loops in the spatial network of length $\ell \geq \tilde{\tau}$, we can identify which atrial regions may act as a substrate for micro-anatomical reentry.

6.4 Materials and methods

In this chapter, we analyse two versions of the spatial network method – a fibre model and a fibre-less null model, shown schematically in Fig. 6.2. Each network is derived from an image-based model of the atria, represented as a cubic mesh of voxels. Vectors in each voxel represent the principal fibre orientation axis at that location.

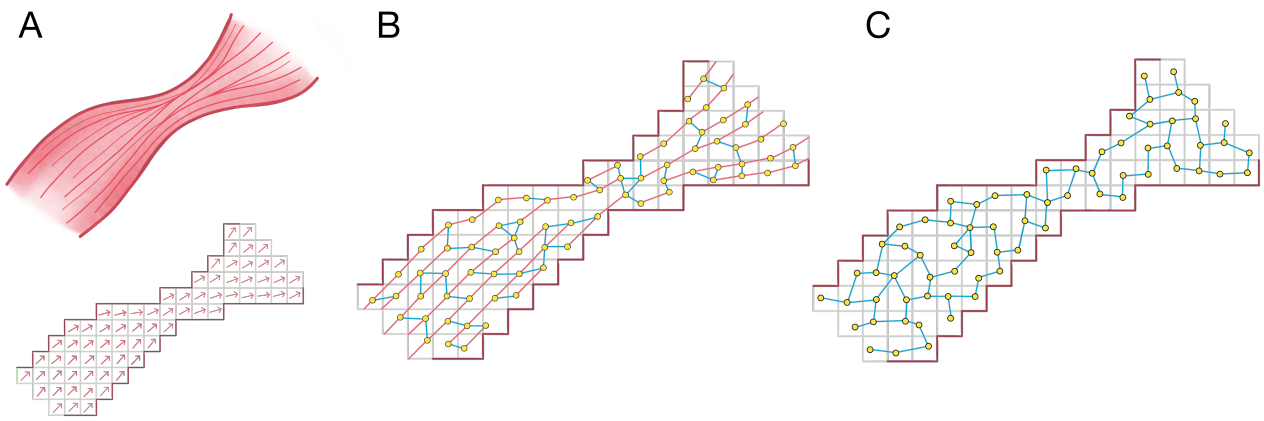

Figure 6.2: (A) An image based model of the atria is derived from atrial imaging (shown in 2d for visual clarity). The model, represented by a mesh of voxels, is coupled with vectors representing the underlying fibre orientation at each point. (B) In the fibre model, a spatial network is generated where nodes (yellow points) are connected with probability 1 along fibre tracts (pink edges), and low probability across fibre tracts (blue edges). (C) In the fibre-less null model, the probability of forming an edge is not biased by the underlying fibre orientation. Figure drawn by Becky Joakim.

In both the fibre and fibre-less null models, nodes are placed throughout the atrial structure at an approximately uniform density of 1 node per voxel. In the fibre model, two nodes are connected by an edge with probability 1 if they are adjacent and in the same fibre tract (pink

edges), and with small probability if they lie in different fibre tracts (blue edges). This introduces a directional bias in the probability that nodes are connected. Conversely, in the fibre-less null model, the probability that two nearby nodes are connected is directionally isotropic; the probability of forming a connection is independent of any underlying fibre structure (blue edges only). Both the fibre and fibre-less null models retain other geometric information such as wall thickness. Hence, differences in the behaviour of the two models can be attributed to the influence of the underlying fibre structure.

6.4.1 Atrial datasets

The spatial network algorithm is tested on three atrial datasets, one human dataset with MRI derived geometry, but synthetic fibre orientations, and two sheep atrial datasets where the fibre structure has been inferred from high resolution serial surface imaging. Each dataset provides a 3d grid of vectors indicating the local fibre orientation at each point in the atria. The sheep datasets are high resolution and inferred directly from the anatomy of each dataset motivating their inclusion here. This lets us assess how the results from our spatial network algorithm differ with different data sources. Summary statistics for the three image-based atrial fibre maps are given in Table 6.1.

Synthetic atrial fibre map

The human dataset is detailed in [196] based on techniques developed in [195]. The atrial geometry used provides no information as to the atrial fibre structure. Therefore, a synthetic fibre map is constructed using a semi-automated rule-based approach.

From a set of 22 user-selected seed points, auxiliary lines and landmarks are drawn using a modified version of the 3d fast marching algorithm [199]. This method constructs fibre bundles by following the shortest path between seed points. Specific fibre structures which do not follow a clear shortest path between seed points, such as the crista terminalis and the pulmonary veins, are accounted for by the inclusion of additional auxiliary lines which form corridors for

	Human	**Sheep Healthy**	**Sheep Heart failure**
Imaging Method	MRI	Serial Surface Imaging	Serial Surface Imaging
Original Image Resolution	$330\mu m$	$50\mu m$	$25\mu m$
Fibre Orientation Extraction	Estimated	Anatomically derived	Anatomically derived
Fibre Orientation Method	Semi-automated rule based approach focusing on key bundles	Eigen-analysis of structure tensor	Eigen-analysis of structure tensor
Resolution after Coarse Graining	$330\mu m$	$300\ \mu m$	$300\mu m$
Dataset Reference	[195, 196]	[197]	[198]

Table 6.1: Summary information for the three image-based atrial fibre orientation datasets.

the fibre bundles along the anatomical feature in question. The inclusion of these auxiliary lines is important since they provide necessary boundary conditions for the fast marching algorithm, see [195] for details.

The method successfully reproduces key fibre bundles; validating the approach, the authors in [195] found that the principal fibre directions showed good correspondence with data and reproduced realistic activation time patterns when the fibre map was used in a standard electrophysiological model. However, a number of artefacts remain at the micro-level. In particular, discretisation effects are found at the tips of the atrial appendages, and the method cannot account for patient specific heterogeneity in the local fibre architecture. The lack of local heterogeneity is likely of particular importance in regions such as the atrial appendages which are known to exhibit disorganised fibre structure in some patients [200].

Although the lack of local heterogeneity cannot be corrected for with this method, discretisation errors are prevented by averaging the local fibre orientation vectors over the five preceding and successive vectors in a given bundle. However, there are a number of small structural holes at

various points in the atria which remain open. The resulting fibre map is significantly smoother than the anatomically derived fibre maps.

Anatomically derived fibre maps

Two anatomically derived sheep atrial datasets are tested using our spatial network method; one healthy sheep acquired in [197] and a sheep with pacing-induced heart failure [198]. These datasets are individual-specific, retaining significant local heterogeneity in the fibre maps, and are provided at high resolution, in contrast to the synthetic human fibre map. In this proof of concept, our aim is not to assess the differences in micro-anatomical reentry between humans and sheep, or between healthy and heart failure (HF) sheep, but to understand how different sources of fibre orientation information affect our results.

The two sheep atrial models were acquired ex-vivo by perfusing the atria at physiologically realistic filling conditions and sealing the atria in parafin wax. The atria were then sliced using an ultramiller in $50\mu m$ ($25\mu m$) steps, and imaged at $8.33\mu m$ ($6.25\mu m$) pixel resolution for the healthy (HF) sheep. Individual images were pre-processed using a suite of techniques to minimise the presence of artefacts. In particular, and in contrast to the human dataset, small holes in the 3d atrial structure were filled and islands of isolated cells were removed from the model. A Gaussian convolution was applied to each image before extracting fibre orientation data. From each image-based model, fibre orientations are estimated from eigen-analysis of the structure tensor, where the principal fibre direction is taken as the eigenvector with the smallest eigenvalue. A second phase of smoothing was applied to the model after estimating the fibre vectors. Although the exact level of smoothing is not specified in [197], local variation in fibre orientation is significantly higher than for the synthetic human datasets.

To validate the structure tensor approach, the authors of [197] and [198] systematically checked the calculated fibre orientations against high resolution images of the individually imaged slices where the orientation of individual cardiomyoctes can be seen. Testing the fibre maps using a simple electrophysiology model, the authors find preferential conduction along key fibre bundles including the crista terminalis, pectinate muscles and the Bachmann bundle.

6.4.2 Spatial network with fibre structure

The schematic shown in Fig. 6.3 outlines the process by which each image-based model is converted into a spatial network in which the underlying fibre structure is retained.

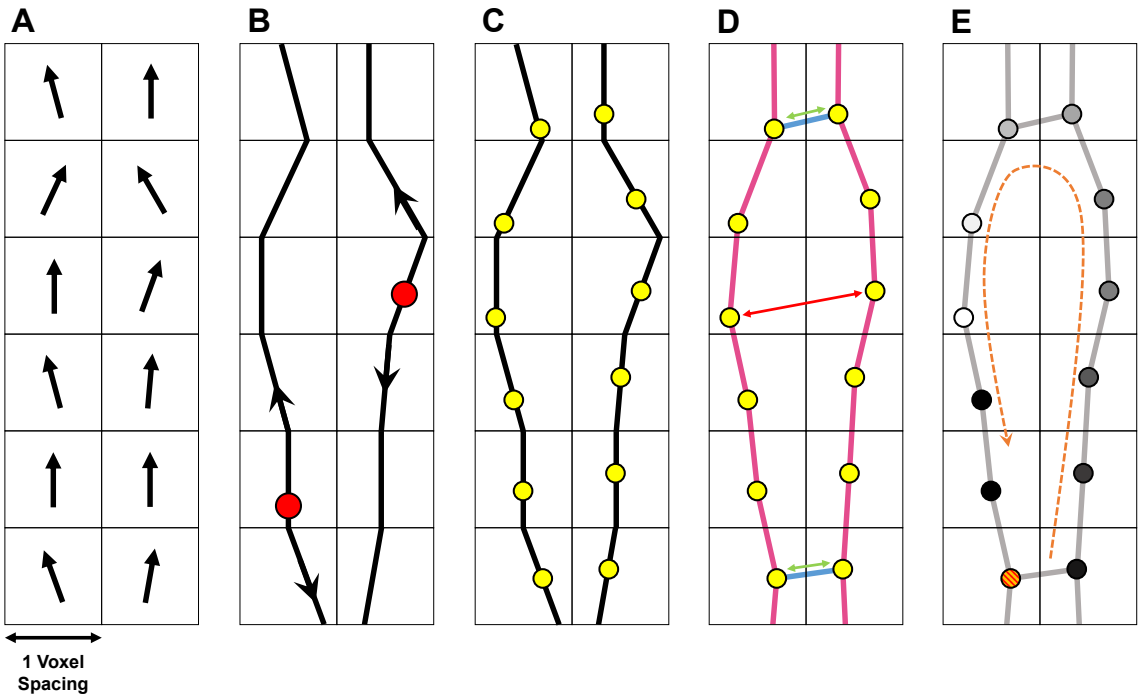

Figure 6.3: A 2d schematic showing the conversion of local fibre orientation vectors into global fibre tracts which are seeded with nodes and coupled into a spatial network. (A) Each image-based model provides a vector field of the local fibre orientation. (B) Beginning from a set of random seed points (red circles), fibres are generated using tractography, propagating forward and backwards in the direction of the local orientation vector. (C) Nodes (yellow circles) are placed at approximately equal intervals along the fibres. (D) Nodes are coupled to neighbouring nodes with probability 1 if nodes are adjacent and in the same fibre (pink edges), and according to a monotonically decreasing function of distance, $p(x; r, c)$, otherwise (blue edges). Nodes which are close together (green arrows) are more likely to be connected than those which are far apart (red arrows). (E) By applying a discrete diffusion model to the network, we identify the micro-reentrant substrate by observing where conduction block (red/orange circle) results in a wavefront (white circle) re-entering a blocked fibre.

Fibre tractography

Fibre tracts are generated across the atrial geometry by applying a modified version of the Evenly Spaced Streamlines (ESS) algorithm, see [201]: (1) Local fibre orientation data for a

given atrial dataset is coarse grained to a specific voxel resolution, Fig. 6.3(A). (2) Placing seed points randomly in the atrial mesh, separated by at least $d_{sep} = 0.7$, global fibre tracts are constructed by propagating forwards and backwards along the local fibre orientation vectors until terminated according to a set of pre-specified rules, Fig. 6.3(B). (3) After fibre tracts have been generated to cover the atrial mesh at an approximately fixed density of 1 node per voxel, nodes are placed evenly along each fibre, Fig. 6.3(C). Importantly, unlike other common tractography methods, this modified ESS approach ensures that the density of fibres is approximately uniform across the atrial structure. Technical details for the tractography process are provided in section 6.9.

Generation of spatial network structure

The fibres generated from tractography are not connected to each other, Fig. 6.3(C). To connect fibres in a single spatial network, nodes are placed at even intervals (steps of 1 voxel length) along fibres starting at a random seed point. These nodes connect to neighbouring nodes in the same fibre with probability $p = 1$. Any two nodes which do not lie on the same fibre are connected with probability

$$p(x; r, c) = \frac{1}{e^{r(x-c)} + 1}, \tag{6.1}$$

where x is the distance between the nodes, $r = 7$ is an arbitrary steepness parameter, and c is a characteristic distance. Connecting nodes according to a distance dependent attachment function is a standard technique when studying percolation on spatial networks [192]. For better performance, only nodes separated by $x < 2$ are considered for connection. By varying the spatial coupling parameter c, the number of transverse connections between fibres can be controlled from a state where fibres are strongly connected at $c \approx 1$, to a state where fibres become increasingly disconnected as $c \to 0$. We hypothesise that this process can effectively simulate the accumulation of electrically insulating interstitial fibrosis in the atria.

It is important to note that Eq. (6.1) is applied equally to all areas of the atria; there is no macroscopic spatial variation in the probability distribution. However, for individual instances of the spatial network, there will be natural fluctuations in the density of connections with

some regions exhibiting a low connection density, and other regions with a higher connection density. Since our datasets do not include individual-specific information on the distribution of fibrosis, we do not augment Eq. (6.1) to account for macroscopic variations in the distribution of fibrosis across the atria in the current work. However, this must be considered in any future work which aims to achieve clinical applicability.

For each dataset, spatial networks were generated with five values of the characteristic distance, c. The specific values used and corresponding risk parameters (see section 6.4.2), λ, are given in Table 6.2.

	Human				
Characteristic Distance, c	0.5250 (L)	0.5000	0.4750 (M)	0.4500	0.4250 (H)
Risk Parameter, λ	0.01569 (L)	0.0190	0.0294 (M)	0.0689	0.2456 (H)

	Sheep: Healthy				
Characteristic Distance, c	0.5500 (L)	0.5250	0.5000	0.4750 (M)	0.4625 (H)
Risk Parameter, λ	0.01900 (L)	0.0250	0.0292	0.0600 (M)	0.1464 (H)

	Sheep: Heart failure				
Characteristic Distance, c	0.5250 (L)	0.5000	0.4750 (M)	0.4625	0.4500 (H)
Risk Parameter, λ	0.0082 (L)	0.0138	0.0410 (M)	0.0754	0.1881 (H)

Table 6.2: Fibre model characteristic distance values, c, and corresponding risk parameter, λ, for each atrial dataset. The labels (L), (M) and (H) indicate the parameter values corresponding to the low, medium and high rate of micro-reentry (risk) spatial networks referred to throughout the text.

Specific parameter choices are labelled as having low, medium or high risk of micro-reentry. These are qualitative markers and are introduced to simplify comparison. The low risk parameter values are chosen by finding the largest characteristic coupling distance for which the risk of micro-reentry is non-negligible; here the rate of micro-reentry, λ, is very low. Different values of the characteristic coupling used across the three datasets reflects differences in the largest characteristic coupling at which micro-reentry is detected. Reductions in the characteristic coupling distance increase the risk of micro-reentry, with the medium risk label being assigned

to intermediate parameter values where the rate of micro-reentry is moderate, and the high risk label being assigned to parameter values where micro-reentry is detected almost immediately after initiating a simulation.

Identifying the substrate for micro-anatomical reentry

The hypothesis we are testing is that the substrate for micro-anatomical reentry equates to regions in the corresponding spatial network where closed loops exist of length $\ell \geq \tilde{\tau}$, where $\tilde{\tau}$ is a typical refractory wavelength. Finding closed loops of a length ℓ in a network is related to finding the set of longest paths between nodes. Unfortunately, this is an NP-hard problem (non-deterministic polynomial-time hardness), such that there is no simple route by which every closed loop can be found [202, Chapter NP-Completeness]. Since we care more about where circuits emerge than the precise path of each circuit, we take an alternative statistical approach whereby individual circuits are identified by applying a discrete diffusion model (DDM) to the network. Other fast approaches which may be appropriate in this context include Eikonal methods, although these would need to be adapted to include an appropriate refractory period [203, 204]. Such an approach should be considered in future work.

The DDM is related to several extremely simple physics models of micro-anatomical reentry [3, 4, 8, 145, 205, 206]. These models are not electrophysiologically realistic models of AF; they do not consider a number of crucial features used in traditional models where ion channel dynamics are simulated with either the mono- or bidomain models describing the propagation of action potentials across the myocardium. However, the strength of discrete models lies in the ease with which structural discontinuities can be modelled, in contrast to conventional models which are significantly complicated when no-flux boundary conditions are imposed at the local level [164]. This justifies the use of the DDM in the current context where we focus exclusively on the structural basis for micro-reentry. However, this approach is not suitable for studying the dynamics of micro-anatomical reentry (for instance how specific drugs may affect micro-reentry). For a wider discussion of the role of model choice when studying structural heterogeneity and the benefits of discrete models see [172].

The rules of the DDM are summarised as follows. First, nodes around the sino-atrial node in the right atrium are labelled and excited every $T = 1000$ discrete time steps. One time step corresponds to the time needed for an excitation to pass from one node to any of its direct neighbours. All nodes which are in the excited state (white nodes in Fig 6.3(E)) at time t_i enter the refractory state (greyscale nodes in Fig 6.3(E)) at $t_i + 1$, during which nodes cannot be re-excited, remaining in this state for $\tau = 83$ time steps. Since the model is parameterised such that wavefronts propagate across 1 edge spacing per time step, this refractory period corresponds to a refractory wavelength of $\tilde{\tau} = 1 \times \tau = 83$ voxel spacings. Converted into real units, $\tilde{\tau} = 83 \approx 25mm$ which is slightly shorter than the fibre tracts identified in [144]. This is very short and corresponds to the lower limit of values which may be observed in highly fibrotic regions during rapid pacing, although we note the specific value is largely unimportant since equivalent results are found by counteracting any change in $\tilde{\tau}$ with an appropriate change in the characteristic coupling, c.

At time $t_i + \tau + 1$, refractory nodes enter the resting state (black nodes nodes in Fig 6.3(E)). Nodes in the resting state at time t_i enter the excited state if (1) at $t_i + 1$ they are directly connected to a node which is excited at time t_i, and (2), conduction is not blocked. Nodes which should enter the excited state at time $t_i + 1$ may fail to do so with small probability $\epsilon = 0.0025$ (orange checked node in Fig. 6.3(E)). This simulates potential unidirectional conduction block at the micro-level although in an arbitrary way; realistic simulation of conduction block would require more sophisticated models, but the uniform block assumption is sufficient for our purposes. The value of $\epsilon = 0.0025$ is chosen to be small to avoid excessive wavefront breakup (i.e., wavefronts should remain approximately planar as is expected during sinus rhythm).

By allowing multiple wavefronts to diffuse across the spatial network, the micro-anatomical substrate is identified when the node at the opening of an isolated fibre tract blocks conduction, such that the discrete wavefront can reenter the isolated fibre from the opposing direction, see Fig. 6.3(E). This process is easily identified in the DDM by the repeated activation of nodes neighbouring the blocked node. All the circuits identified in this manner are at least length $\ell \geq \tilde{\tau}$. Further details are provided in section 6.9.2 including additional schematics.

Defining risk of micro-anatomical reentry

To derive a measure for the risk of micro-reentry in each spatial network we return to the DDM and note that all nodes in the network are equally susceptible to conduction block. If each structure has a similar number of nodes which, if blocked, may initiate a micro-anatomical reentry, then the probability of initiating any given micro-reentrant circuit is approximately constant. Hence, for a spatial network with N potential substrates, the rate, λ, at which new micro-reentries are initiated in the DDM is directly related to the number of substrates N. This argument follows directly from similar arguments in [8].

A process where events take place independently at a constant rate λ is known to follow a Poisson point process.[1] Hence, from the start of a DDM simulation, the waiting time until the first micro-anatomical substrate is identified follows an exponential decay distribution, $\rho(t) \propto e^{-\lambda t}$. Fitting these distributions on a log-linear scale lets us extract the rate parameter λ, which varies depending on the choice of characteristic coupling distance, c, in each spatial network. The rate λ is directly related to the number of potential micro-reentrant substrates in each spatial network. Therefore, λ can be used as a measure for the risk of micro-reentry in our framework.

The activation rate λ measures the risk of micro-reentry for each spatial network as a whole. In a local area, we consider a specific voxel, v, and calculate the number of circuits which form in that particular voxel of the DDM, N_v, when sampling a fixed number of identified circuits, N. The rescaled rate of reentry for the voxel v is then given by

$$\tilde{R}_v = \frac{\lambda \cdot N_v}{\sum_v N_v}, \tag{6.2}$$

where the denominator provides a normalisation. When plotting the rescaled rate of reentry onto the atrial surfaces, we smooth $\tilde{R}_v$ using a 3d Gaussian convolution with a standard deviation of 5 voxels ($\sim 1.5mm$).

[1] It is interesting to note that the Poisson point process has recently been used to explain the generation and destruction of phase singularities during fibrillation, see [207, 208].

6.4.3 Spatial network without fibre structure

In this section, we consider a null model where the spatial network excludes fibre orientation. All other geometric information is retained. Hence, regions which are identified as "at risk" in the fibre-less null model emerge due to geometric factors, whereas differences between the fibre and fibre-less models can be attributed to the fibre structure. The fibre-less spatial network is generated analogously to the fibre model, with the omission of the initial fibre tractography steps, illustrated schematically in Fig. 6.4.

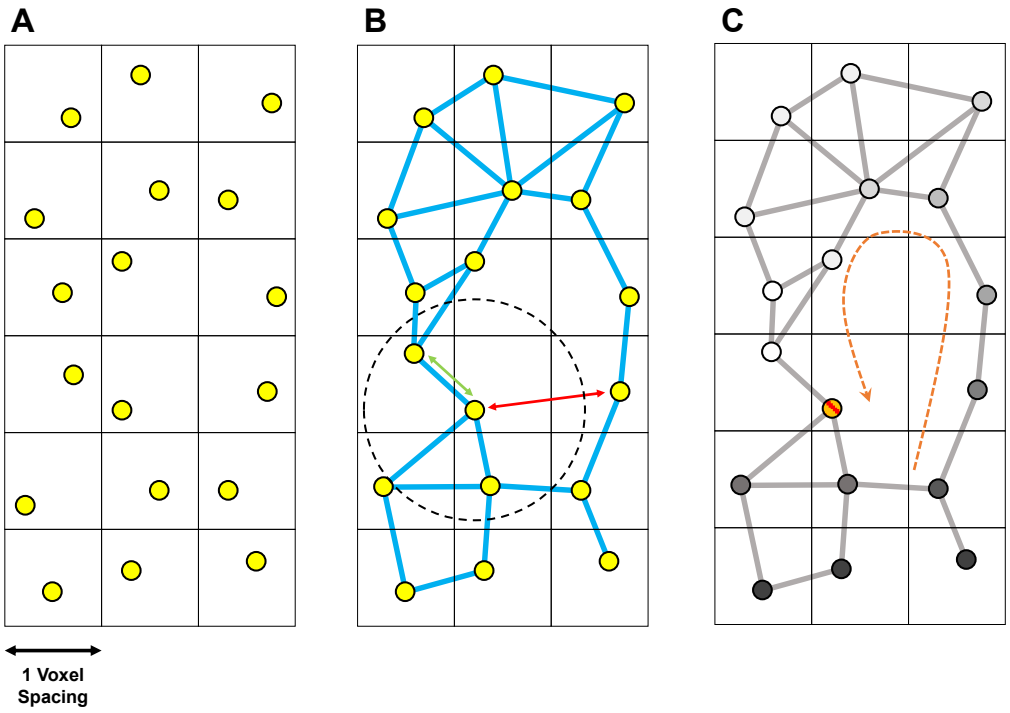

Figure 6.4: A 2d schematic demonstrating the construction of the spatial network using the fibre-less null model. The process is similar to Fig. 6.3 where steps (A) and (B) are skipped. (A) Nodes are distributed uniformly in each voxel. (B) Nodes are connected according to a monotonically decreasing function of distance, $p(x; r, c)$, Eq. (6.1), where c is a characteristic distance (dashed circle). There is no directional bias in the connection probability (blue edges only). Nodes which are closer together than the characteristic distance (green arrows) are more likely to be connected than those which are far apart (red arrows). (C) The micro-reentrant substrate is identified using the same process as the fibre model, see Fig. 6.3.

Nodes are randomly distributed across each atrial mesh at a density of approximately 1 node per voxel and are added to the spatial network if the distance between the new node and any existing nodes is greater than or equal to $d_{sep} = 0.7$; this prevents an over-density of nodes in

any one location.

To construct the spatial network, nodes are connected using the same distance dependent probability function as the fibre-model given in Eq. (6.1). Unlike the fibre-model, none of the nodes are already connected by edges which represent the underlying fibre structure of the atria (no pink edges); the probability of connecting two nodes is independent of their orientation. The regions at risk of micro-reentry are identified using the same DDM described in section (6.4.2).

For each atrial dataset, we generate spatial networks for three values of the characteristic coupling, corresponding to low, medium and high risk cases, see Table 6.3. The characteristic coupling values in Table 6.3 are significantly larger than the coupling values for the fibre model, see Table 6.2, since there are no permanent longitudinal edges.

	Human		
Characteristic Distance, c	0.6500 (L)	0.6375 (M)	0.6250 (H)
Risk Parameter, λ	0.0240 (L)	0.0527 (M)	0.2775 (H)

	Healthy Sheep		
Characteristic Distance, c	0.6375 (L)	0.6250 (M)	0.6125 (H)
Risk Parameter, λ	0.0141 (L)	0.0310 (M)	0.1534 (H)

	HF Sheep		
Characteristic Distance, c	0.6500 (L)	0.6250 (M)	0.6175 (H)
Risk Parameter, λ	0.0161 (L)	0.0520 (M)	0.2152 (H)

Table 6.3: Fibre-less null model characteristic distance values, c, and corresponding risk parameter, λ, for each atrial dataset. The labels (L), (M) and (H) indicate the parameter values corresponding to the low, medium and high rate of micro-reentry (risk) spatial networks referred to throughout the text.

.

6.4.4 Analysis methods

Fibre orientation correlation & longitudinal connection fraction

In the current chapter, we have focused exclusively on generating spatial networks from pre-existing image based models of the atria which include fibre orientation data. Therefore, we must treat the pre-existing fibre datasets as a ground truth and validate that our spatial network methods accurately capture the properties of the underlying fibre dataset. The individual datasets have been validated previously, see section 6.4.1.

Each fibre map is defined by a vector field, where the vector at each coordinate corresponds to the estimated fibre direction in that specific voxel. We can measure the local heterogeneity in the fibre map by calculating the fibre orientation correlation (FOC), defined as the average value of the absolute dot product of each voxel's fibre orientation vector with the fibre orientation vector of the voxel's six immediate neighbours (on a 3d cubic mesh). Neighbours outside the atrial geometry are excluded from the calculation. Taking the absolute value ensures that the FOC does not distinguish between parallel and anti-parallel vectors. Each of the vectors is normalised such that the FOC can take a minimum value of 0, corresponding to a vector which is orthogonal to all of its neighbours, and a maximum value of 1, where all vectors are aligned.

We would like to extract an equivalent measure from each spatial network for validation purposes. To do this, we calculate the longitudinal connection fraction (LCF), defined as the number of longitudinal edges (pink edges in Fig. 6.2) attached to each node in a voxel, normalised by the total number of edges connected to each node, either longitudinal or transversal (blue edges in Fig. 6.2). This measure is chosen as a proxy for fibre alignment since there is no natural concept of fibre alignment in the spatial networks. In regions with a high density of aligned fibre tracts (as generated using tractography, see Fig. 6.3(C), we expect the LCF to be large, whereas in regions with disordered fibres, we expect the LCF to be smaller.

Although geometric factors also impact the LCF, such as proximity to the atrial walls, we expect this effect to be small. If the spatial networks accurately capture the underlying heterogeneity of each fibre dataset, we expect to find a linear relation between the LCF and the FOC.

Spatial clustering of reentrant circuits

For each spatial network, we measure how spatially clustered the identified micro-reentrant substrate is. To do so, we sample 1000 circuits from the set of identified risk substrates for each network and extract the coordinate of each detected circuit. The number of independent spatial clusters is then derived by using the DBSCAN algorithm, a standard algorithm for spatial data clustering, see [209]. We set `eps` $= 10$ ($\approx 3mm$), which is the parameter controlling the distance between points for these to be considered in the same cluster. The number of points required for a cluster is set to 1. If only one cluster is identified by DBSCAN, this indicates that all 1000 sampled circuits in a given network fall within a single, small region of the atrial structure. Conversely, if 1000 clusters are identified, no two circuits are identified in the same location.

One possible interpretation of the clustering is that if all micro-reentrant circuits are found in the same region, these will be easy to destroy or isolate via catheter ablation. However, if circuits are widely distributed, isolating all of these circuits will be difficult, if not practically impossible.

Atrial wall thickness & occupied voxel fraction

Wall thickness is a common measurement used when analysing atrial geometry. However, in the bulk many perpendicular surface vectors pass through the same voxel so that the thickness is not uniquely defined.[2] To resolve this issue, we introduce a novel measure analogous to wall thickness, the local occupied voxel fraction (OVF), which assesses the proximity of voxels to the atrial walls and captures differences in the local wall curvature and thickness gradients, see Fig. 6.5.

The OVF is defined as the number of voxels within a radius $\tilde{r} = 5$ (approximately 1.5mm) which are inside the atrial structure, normalised by the total number of voxels within a sphere of radius $\tilde{r}$. The OVF is not normalised by wall thickness, capturing the qualitative properties

[2]An illustrative example is to consider an ellipse. Both the major and minor axes of an ellipse are straight lines, perpendicular to the surface, passing through the ellipse centre. Although the thickness of the ellipse can be defined at each point on the ellipse surface, at the ellipse centre, both the length of the minor and major axes could be interpreted as the thickness.

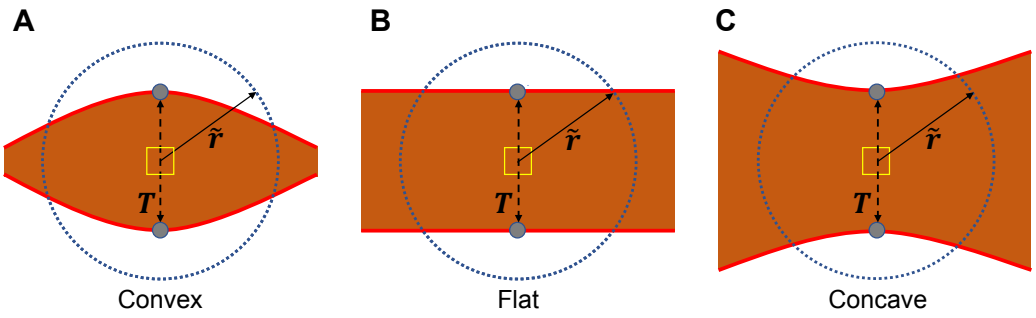

Figure 6.5: A 2d schematic illustrating the definition of the occupied voxel fraction (OVF); an analogue to wall thickness which is well defined in the atrial bulk. Consider three regions with equal thickness T. If we measure the fraction of myocardial voxels (brown region) enclosed within a sphere of radius $\tilde{r}$ (dotted blue circle), relative to the volume of the sphere, we note that the occupied voxel fraction is lowest for convex wall morphologies, moderate for flat morphologies, and largest if the walls are concave. Increasing (decreasing) the thickness T typically increases (decreases) the corresponding OVF values.

of traditional wall thickness measurements implicitly. The OVF increases (decreases) if the atrial walls become thicker (thinner), and distinguishes between regions with convex, flat, or concave wall morphologies. A natural consequence of this method is that branchy regions of the atrial myocardium, such as the pectinate muscles, are identified as having a low OVF, whereas voxels in the atrial bulk which are more than 1.5mm from the surfaces will have an OVF of 1.

Qualitatively comparing the OVF and thickness, see section 6.9.3, we note that the OVF appears smoother than the thickness measurement, and is subject to fewer anomalous outlying values (such as very large thickness measurements at the end of a vessel opening). The OVF is mathematically simple and less computationally expensive than traditional thickness measurements.

6.5 Results

6.5.1 Spatial distribution of identified micro-reentrant substrate

Figures 6.6 & 6.7 show the substrate identified as susceptible to micro-reentry in the human atria using the fibre and fibre-less null models respectively, at low, medium and high risk. The risk level refers to the overall risk of micro-reentry across the whole atrial tissue, not to the risk

in a specific region. The risk distributions for the sheep atria are shown in section 6.10.

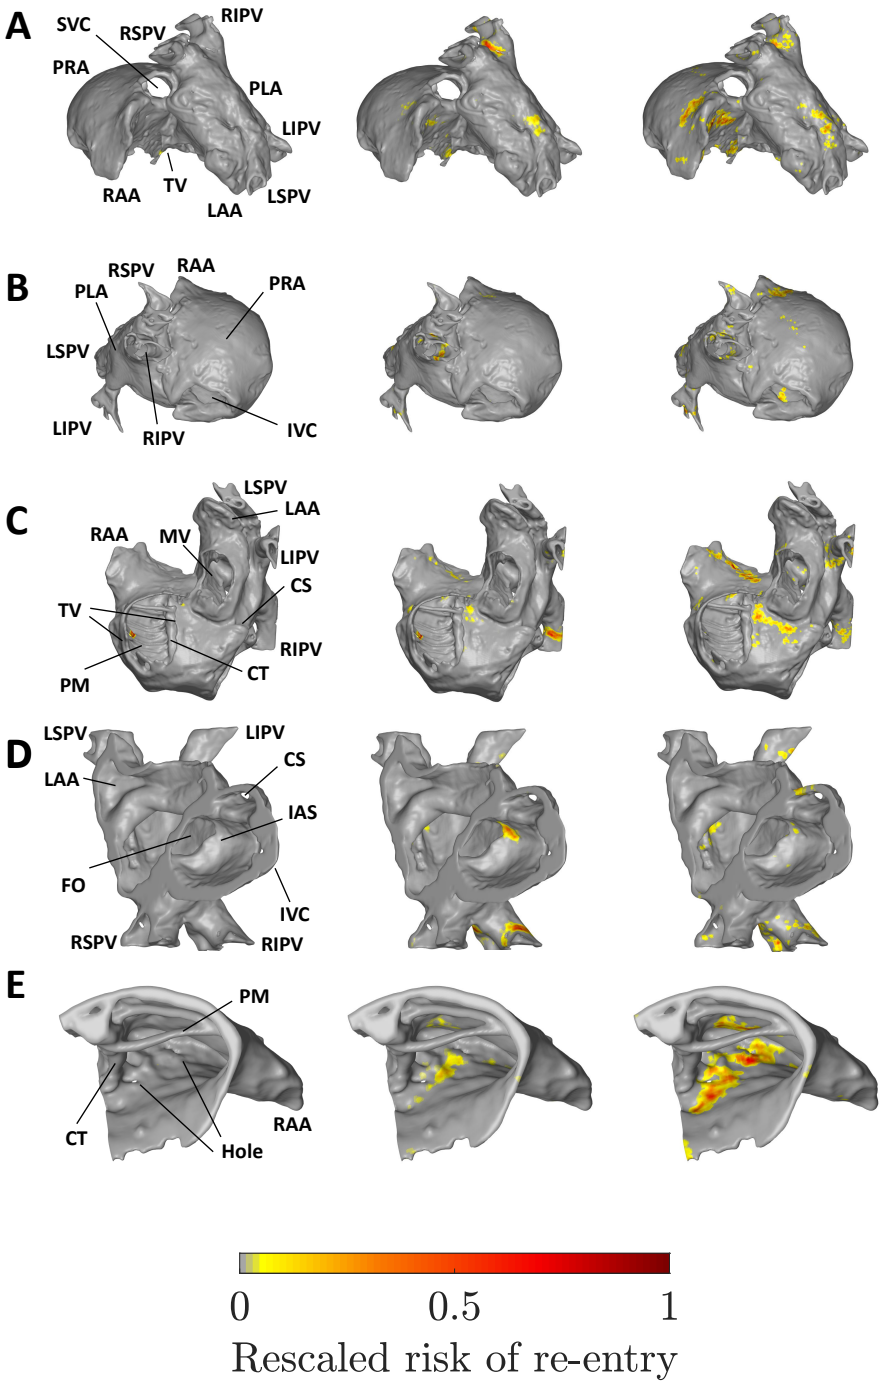

Figure 6.6: The micro-reentrant substrate for the human atria at low (left column), medium (middle) and high risk of micro-reentry (right). A: Superior view. B: Posterior view. C: Inferior view. D: Cut-through. E: Superior right atrial zoom. IVC/SVC: Inferior/ superior vena cava. PRA/PLA: Posterior right/ left atrium. RAA/LAA: Right/ left atrial appendage. TV/MV: Tricuspid/ mitral valve opening. CS: Coronary sinus. FO: Fossa ovalis. PM: Pectinate muscles. CT: Crista terminalis. IAS: Inter-atrial septum. RIPV/LIPV/RSPV/LSPV: Right/ left, inferior/ superior, pulmonary vein.

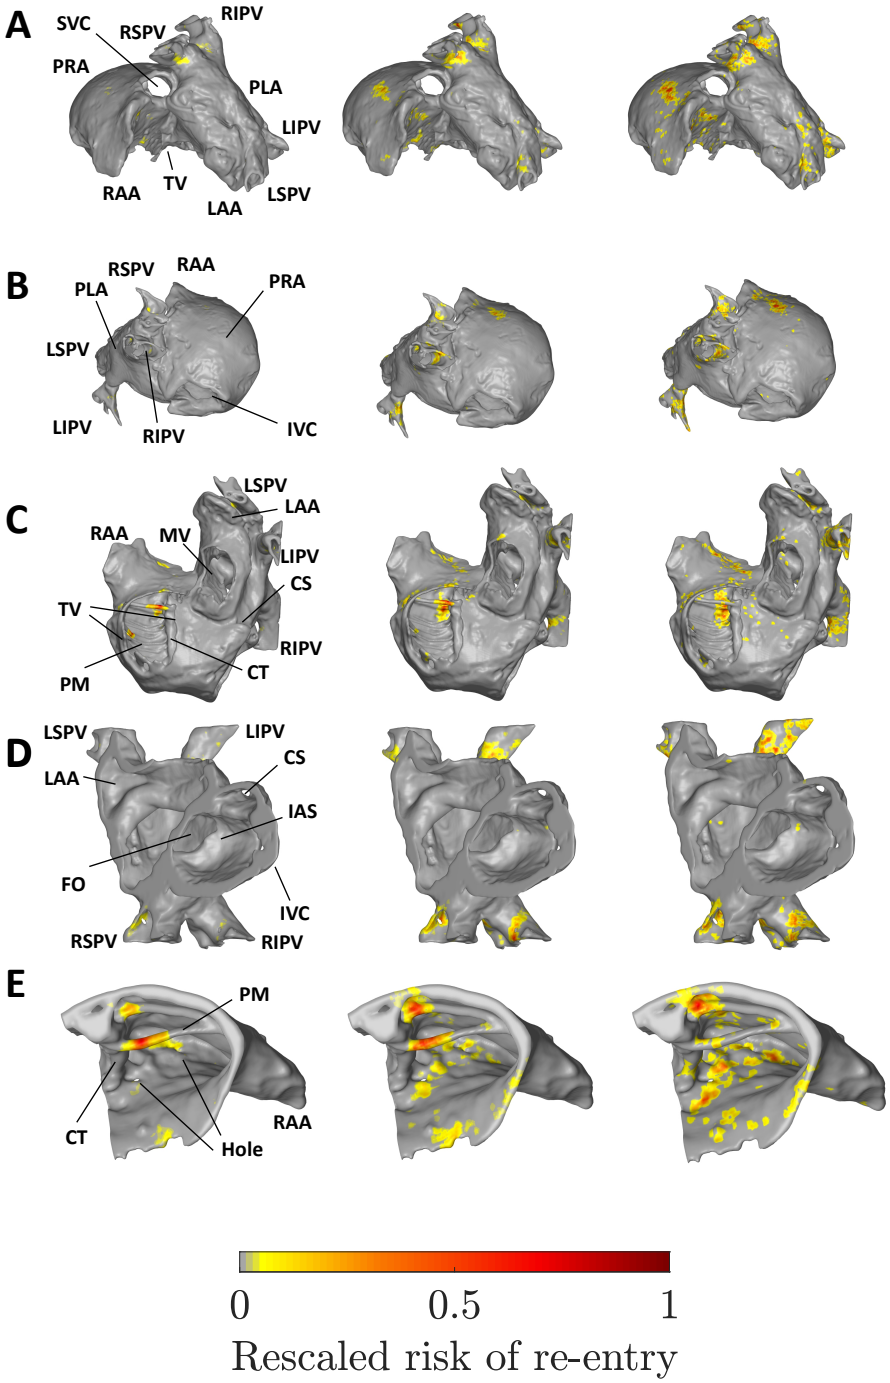

Figure 6.7: The micro-reentrant substrate for the human atria using the fibre-less null model at low (left column), medium (middle) and high risk of micro-reentry (right). A: Superior view. B: Posterior view. C: Inferior view. D: Cut-through. E: Superior right atrial zoom. See Fig. 6.6 for labels.

The reader should be conscious that the risk distributions shown do not represent the regions at risk for a single atrial tissue, but rather show the collation of the risk regions identified across 1000 randomly generated spatial networks. If a region is highlighted as at risk in the left column, this may be interpreted as an area which is susceptible to micro-reentry even if the fibrosis density is low (but non-zero). Regions identified in the middle and right columns are regions susceptible to micro-reentry at moderate and high fibrosis densities. This illustrates our key result: different regions of the atria are susceptible to micro-reentry at different fibrosis densities. In particular, regions may have a characteristic fibrosis density range within which they are susceptible to micro-reentry; if the fibrosis density is too high or low these regions will not exhibit micro-reentry. This result is related to the observation that reentry often anchors to fibrotic border zones, rather than to regions with zero fibrosis, or to regions with very compact fibrosis [210].

Micro-reentrant substrate in spatial network with fibre structure

At low risk, micro-reentry is exclusively observed along a spatially isolated ridge of the pectinate muscles (PMs) in the right atrium. At medium risk, the pulmonary veins (PVs), in particular the right inferior pulmonary vein (RIPV), emerge as the dominant micro-reentrant substrate. Secondary substrates include the junction of the left LSPV and the PLA, and the inter-atrial septum (IAS). However, the hole in the IAS may play a role in it emerging as a risk region. An additional strip of low risk is observed in the superior right atrium, adjacent to the crista terminalis (CT), see Fig.6.6(E). The strip in question lies between two small holes in the atrial dataset, with sub-millimetre diameters.

At high risk, the susceptible substrate is significantly more diffuse than at lower risk. However, there remain multiple regions with very low or zero risk, in particular the posterior right atrium (PRA), left atrial appendage (LAA) and inferior vena cava (IVC). The dominant risk substrate lies along the strip between two small holes in the superior right atrium. Secondary regions of high risk are maintained in the PVs, with risk migrating slightly further from the LSPV junction into the PLA. The opening of the coronary sinus forms a new substrate which is not

observed at higher coupling. The risk observed previously at the IAS and PMs is largely absent.

Micro-reentrant substrate in spatial network without fibre structure

The risk substrate identified by the fibre-less null model is qualitatively similar to the risk regions identified in the fibre model. Key risk regions including the superior right atrium and the PV sleeves are common to both models. Likewise, the PRA, LAA and IVC are not susceptible to micro-reentry in either model.

At low risk, the dominant risk substrate in the fibre-less null model is confined to two spatially isolated PM ridges, one of which exhibits risk in the fibre model. Additional diffuse risk is observed in the superior right atrium and in the right PV sleeves. At medium risk, the low risk substrate is retained, with additional risk in all four PVs and diffuse risk across the superior right atrium. These risk regions are consolidated at high risk, with a small reduction in risk observed along the spatially isolated PV ridges. The substrates observed for the fibre model along the coronary sinus (CS) opening and at the IAS are not observed in the fibre-less null model.

6.5.2 Spatial clustering of micro-reentrant substrate

The spatial distributions of the micro-reentrant risk substrate for the fibre and fibre-less null models suggest that the risk substrate is spatially confined for both models at low risk, but that the substrate covers a wider area as the risk of micro-reentry increases. There is evidence that the substrate is more clustered for the fibre model than for the fibre-less null model.

We quantify the number of independent spatial risk clusters using the DBSCAN algorithm from 1000 randomly sampled micro-reentrant circuits in each spatial network, shown in Figs. 6.8(A), (B) and (C) for the human, healthy sheep, and HF sheep atria respectively. Since the number of sampled circuits is fixed, we note that the level of clustering is the inverse of the number of clusters, such that a small number of clusters is equivalent to high clustering, and vice versa.

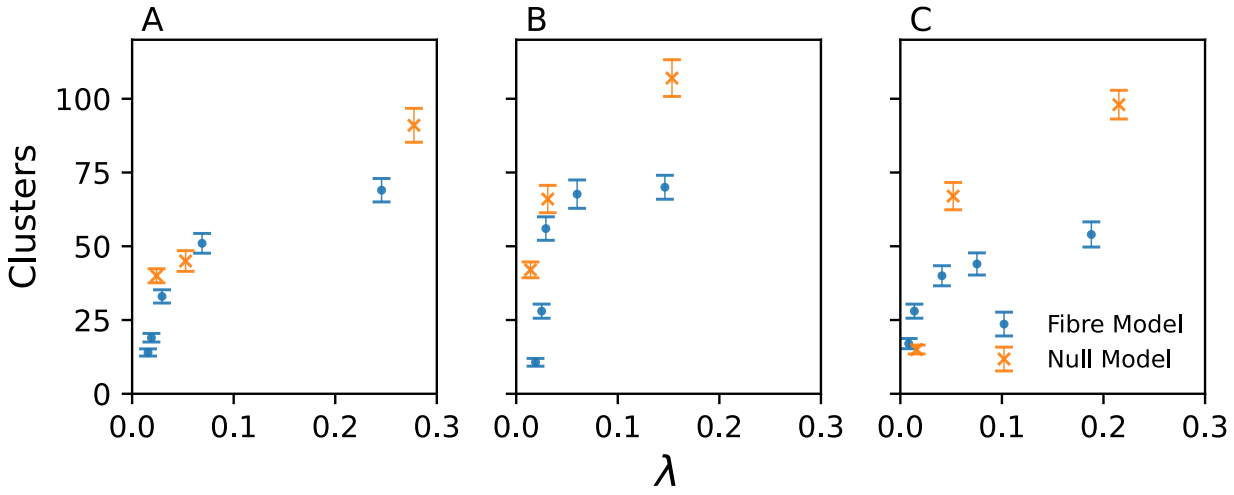

Figure 6.8: The number of independent spatial clusters identified as susceptible to micro-anatomical reentry in the fibre model (blue points) and the fibre-less null model (orange crosses) for (A) the human atria, (B) the healthy sheep atria, and (C) the HF sheep atria, as a function of the risk parameter, λ. Each clustering value is calculated from 1000 sampled micro-reentrant circuits. Error bars from bootstrapping.

At low risk (small λ) the number of clusters is small for both the fibre and fibre-less null models indicating that most of the identified micro-reentrant circuits are confined to a small number of regions. Conversely, as the risk of micro-reentry increases, the number of clusters grows rapidly for both models. At the highest risk values, the substrate in the fibre-less null models is spread across a significantly larger number of clusters than the equivalent fibre models, suggesting that fibre structure plays a key role in spatially confining the micro-reentrant risk.

6.5.3 Occupied voxel fraction

Figure 6.9 shows the average occupied voxel fraction for the identified micro-reentrant substrates in the fibre and fibre-less null models. For all three atrial datasets, the figure demonstrates that at low risk, the mean OVF is minimised and significantly below the average value for each dataset. As risk increases (increasing λ), the mean OVF increases. This indicates that the bias to thin convex regions of the atria is reduced as edges are progressively removed in each spatial network. Framed in terms of fibrosis densities, this result illustrates that regions with low OVF are susceptible to micro-reentry at a lower fibrosis density than regions with

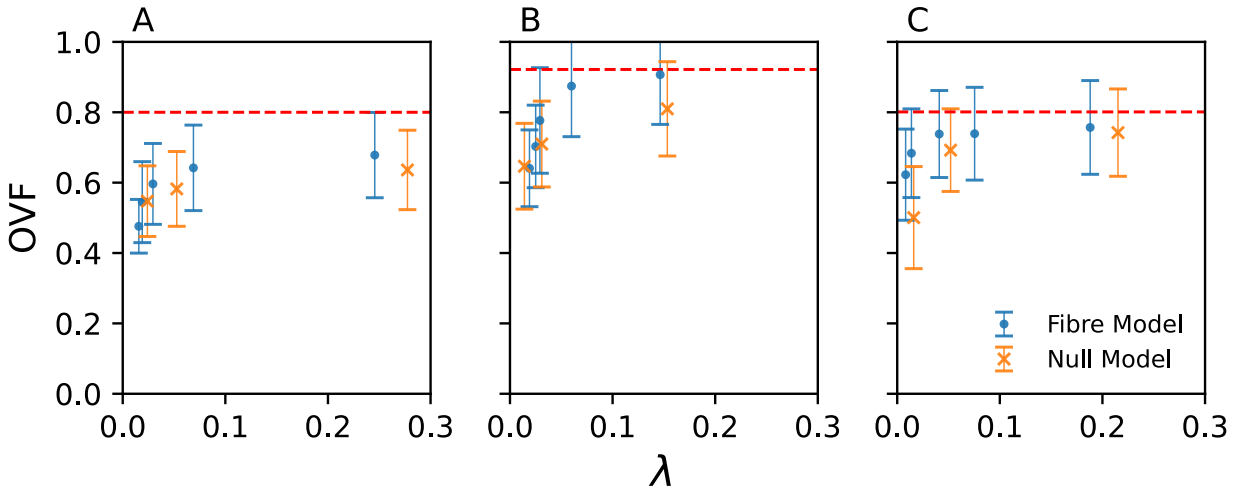

Figure 6.9: The occupied voxel fraction averaged over voxels where micro-reentry is detected in the fibre model (blue points) and the fibre-less null model (orange crosses) for (A) the human atria, (B) the healthy sheep atria, and (C) the HF sheep atria, as a function of the risk parameter, λ. The dashed line in each subfigure is the average value of the occupied voxel fraction for all the voxels in the atrial geometry.

larger OVF. Note that the bias to low OVF at small λ does not guarantee that a specific low OVF region will harbour micro-reentry; other factors influence the risk substrate including the degree of longitudinal coupling in a region.

To check whether the differences in mean OVF values are statistically significant we apply a permutation test [211]. Permutation tests are non-parametric and therefore avoid issues that arise due to the non-normality of the OVF distributions. Comparing the OVF in voxels where micro-reentry is detected to an equivalently sized sample drawn uniformly at random from each geometry, we test the null hypothesis that both samples are drawn from the same distribution. With the exception of the OVF distribution for the healthy sheep at high risk in the fibre model (see Table 6.2), we can reject the null hypothesis with $p < 0.01$, see section 6.10.3 for details on the statistical analysis.

6.5.4 Comparing fibre and fibre-less micro-reentrant substrates

Figure 6.10 shows an example of the micro-reentrant substrate for the fibre and fibre-less null models alongside their corresponding OVF and LCF values in the superior right atrium. The

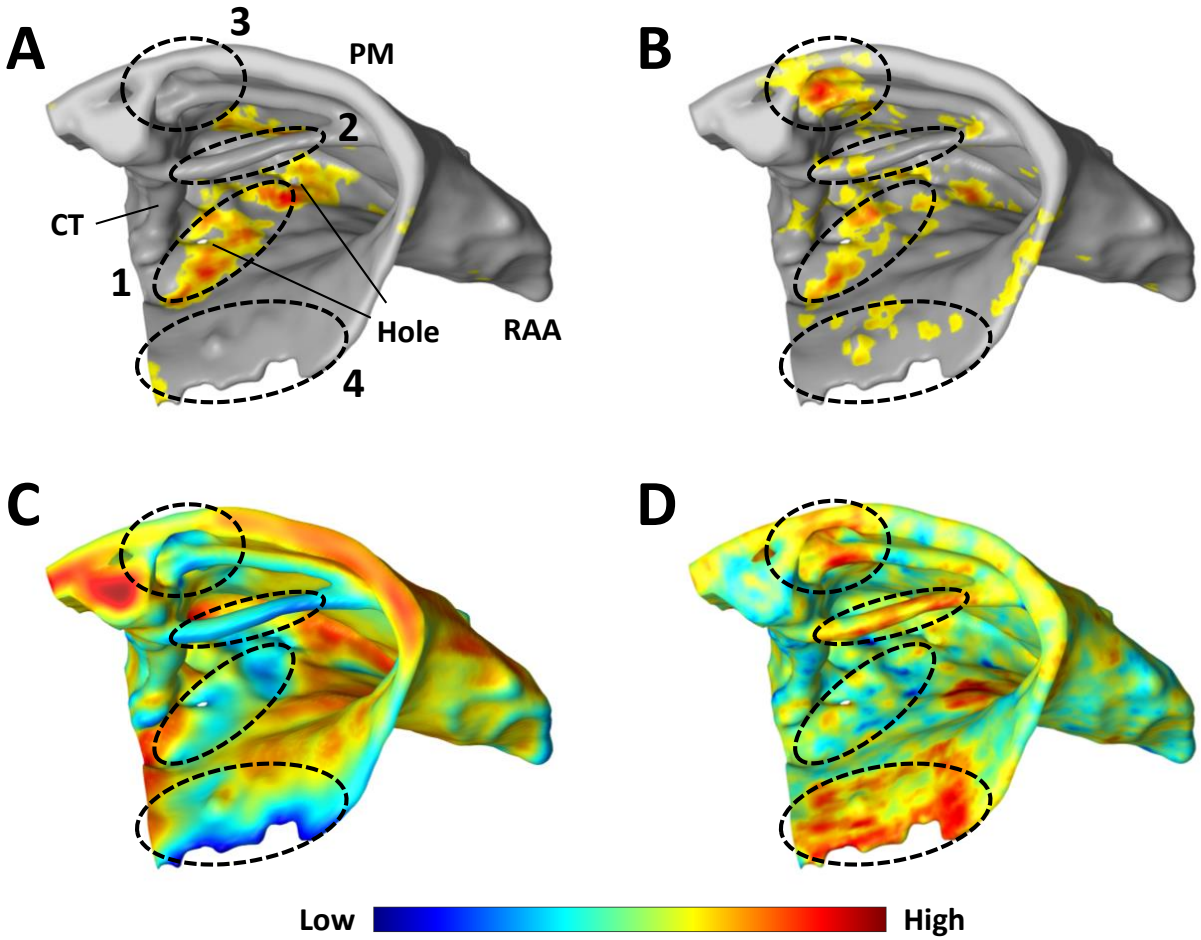

Figure 6.10: A view of the human superior right atrium illustrating factors influencing the spatial distribution of risk in (A) the fibre model, and (B) the fibre-less null model. (C) The occupied voxel fraction (OVF). (D) The longitudinal connection fraction in the fibre model (LCF). Regions (1–4) are referenced in the main text. Colourbar limits: OVF $\in [0.3, 0.9]$, LCF $\in [0.4, 0.8]$.

figures is for the human atria at high risk (right column in Figs. 6.6 and 6.7) where the bias to low OVF is small but non-zero.

In the superior right atrium, the micro-reentrant substrate at high risk is concentrated along a strip between two small structural holes (region 1) for the fibre model, Fig. 6.10(A). This region exhibits moderate risk in the fibre-less null model, Fig. 6.10(B), with additional risk along the spatially isolated PM (region 2), the junction of the CT and a second PM ridge (region 3), and around the opening of the TV (region 4).

Region 2 exhibits low risk in the fibre model, despite being a low OVF region, but is the dominant risk region for the fibre-less null model when the rate of micro-reentry is low or

medium, see Fig. 6.7(E) left and centre columns. This is because the strong longitudinal coupling in region 2 suppresses micro-reentry in the fibre model, but this effect is absent in the fibre-less null model. This also illustrates that region 2 is most susceptible to micro-reentry in the fibre-less null model at low and moderate fibrosis densities, but that the risk reduces at very large fibrosis density, in contrast to some thicker regions of the atria. Regions 3 and 4 show very low or zero risk in the fibre-model.

Comparing the identified risk regions to the OVF values shown in Fig. 6.10(C), we note that regions 1–4 all exhibit low OVF values (blue on the colour scale). Conversely, other regions in the superior right atrium are at larger OVF and do not exhibit risk in either the fibre or fibre-less null models.

Recall that the LCF is only defined for the fibre model; in the fibre-less null model the probability of connecting two nodes is independent of orientation. Of the four regions identified in Fig. 6.10(C) with low OVF values, we observe that the LCF values, Fig. 6.10(D), in region 1 are low (blue), whereas strong longitudinal coupling is observed in regions 2–4. This suggests that the risk of micro-reentry is suppressed in regions with high LCF, but can be enhanced in regions with low LCF.

To quantify these patterns, we plot the distribution of LCF and OVF values across all the voxels in the human atria. Figures 6.11(A) & (D) show the distribution where all voxels are equally weighted at (A) low and (D) high micro-reentrant risk. In Figs. 6.11(B) & (E) these distributions are weighted by the rescaled risk of micro-reentry observed in each voxel for the fibre model. The equivalent is shown in Figs. 6.11(C) & (F) where voxels are weighted by their risk in the fibre-less null model. Although the LCF is not defined for the fibre-less null model, it is illustrative to label null model voxels according to the corresponding fibre model LCF values.

Comparing Figs. 6.11(A) & (D), we note that the distribution of OVF values is equivalent at low and high risk, whereas the distribution of LCF values appear to be uniformly translated to larger LCF. This is expected since the progressive removal of transverse connections in the fibre model increases the LCF, but has no effect on the OVF. The high density strip at OVF ~ 1

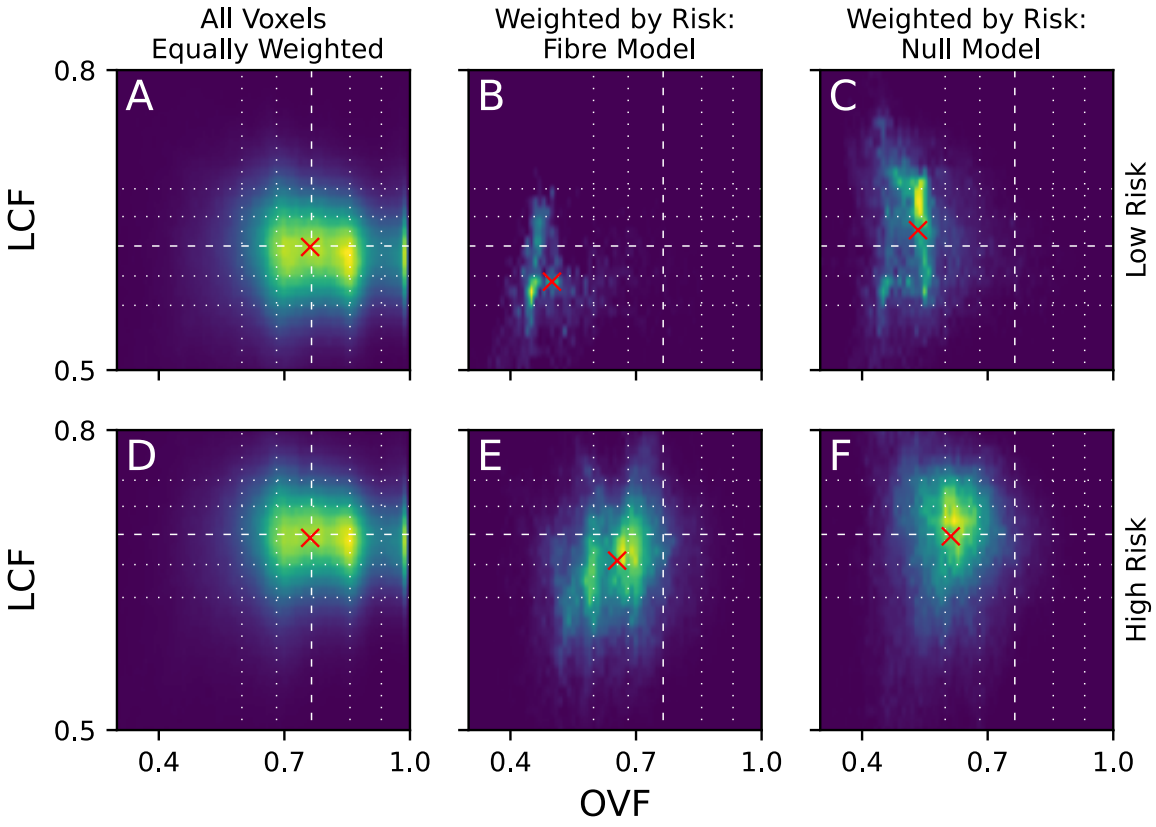

Figure 6.11: The distribution of LCF and OVF values in the human atria at low (A-C) and high (D-F) risk. (A) & (D) The distribution across all the voxels in the human atria. (B) & (E) The distribution weighted by the relative risk of reentry in each voxel for the fibre model. (C) & (F) The equivalent for the fibre-less null model. Note, the LCF is not defined for the fibre-less null model, but for illustrative purposes we include the LCF for the equivalent fibre model. Dashed lines: median. Dotted lines: 10^{th}, 25^{th}, 75^{th} and 90^{th} percentiles. Red crosses: weighted mean.

corresponds to bulk voxels in regions where the atrial walls are thick.

If the distribution of micro-reentrant risk is unbiased by OVF or LCF, the distributions in Figs. 6.11(B), (C), (E) & (F) should follow the same patterns shown in Figs. 6.11(A) & (D). For the fibre model at low risk, Fig. 6.11(B), the voxel distribution is significantly skewed for both the OVF and the LCF with the weighted mean falling in the first OVF decile and the first LCF quartile. For the fibre-less null model at low risk, Fig. 6.11(C), the mean OVF value also falls in the first decile, but there is a small skew to LCF values larger than the median. This is an indirect effect and is explained by the strong longitudinal coupling along spatially isolated regions such as the PMs.

For both the fibre and fibre-less null models it is interesting to note that the risk density lies

almost exclusively in the first OVF quartile, with negligible density at higher OVF values. In contrast, LCF values are clearly skewed in the fibre model, although some risk density remains at larger LCF values. This implies that at low risk, the OVF is the dominant factor in determining the micro-reentrant risk substrate, with fibre structure playing an important secondary role.

As the risk of micro-reentry increases, we observe that the bias to low OVF values remains, but is slightly reduced relative to the low risk cases, see Figs. 6.11(C) & (F). However, the risk of micro-reentry is still completely suppressed at the largest OVF values. For the fibre model, a small bias towards reduced LCF is observed. In contrast, the fibre-less null model shows minimal bias in the distribution of LCF values. Interestingly, the increased dependence on the LCF in the fibre model seems to slightly reduce the bias towards the low OVF values observed in the fibre-less null model, although the effect appears small. Statistical arguments supporting the above discussion are provided in section 6.10.3.

6.5.5 Comparing atrial datasets

The distributions of micro-reentrant risk for the anatomically derived sheep atria are shown in section 6.10. Qualitatively, our results are robust across all three datasets. At low risk, the micro-reentrant substrate is highly clustered in a small number of regions. These regions correspond to areas of low OVF. In the fibre-less null models, this bias reduces with increasing risk. In particular, the probability of observing micro-reentry in the null model at OVF ~ 1 is approximately zero across all spatial networks. In the fibre models, a bias to low OVF is observed, but the effect is smaller than for the fibre-less null models.

In the fibre models, there is a bias at all risk levels to areas of reduced longitudinal coupling (LCF). However, effect sizes are significantly larger for the synthetic human atrial dataset than for either sheep atrial dataset, and the effect is less influential than the bias to low OVF. For the HF sheep, at medium and low risk we find a small bias to areas of increased LCF, although visual inspection suggests that these locations correspond to regions with particularly low OVF values. The healthy sheep atria at high risk is the only spatial network in which we observe

any significant micro-reentrant risk in regions with OVF ~ 1. However, the LCF values for the majority of those voxels falls around the 25th percentile, suggesting that only with reduced longitudinal coupling can micro-reentry emerge in thick regions of the atrial bulk.

These results suggest that the spatial network approach is relatively robust for datasets with both synthetic and anatomically derived fibre structures. Despite this, the robustness of the method cannot correct for the intrinsic limitations of individual datasets. In particular, there are two key differences between the synthetic human dataset and the anatomically derived sheep datasets which appear to affect our results.

Firstly, the local fibre heterogeneity in the human atria is not realistic, with individual heterogeneities likely arising at the boundaries of the user-defined fibre bundles, see section 6.4.1. As a consequence, over 90% of the voxels in the datasets are in regions with very high fibre alignment ($FOC > 0.9$), artificially suppressing micro-reentry across vast regions of the atria, and promoting micro-reentry artificially along fibre bundle boundaries. In this proof of concept, these limitations are acceptable for testing our methods, but the results derived from the fibre-model cannot accurately reflect how a real human fibre structure may affect the substrate for micro-reentry. Although both sheep atrial datasets have undergone some smoothing, heterogeneities in these datasets have been histologically validated and are distributed with finer detail across the atria.

The second difference concerns the pre-processing and resolution of the atrial structures. Although both sheep datasets were down-resolved to a voxel resolution that is comparable to the human resolution, visual inspection suggests that the sheep still retain significantly finer structural detail than the human atria. In particular, the fine muscular ridges of the PMs (with low OVF) are preserved in the RAA and LAA in the sheep atria, whereas these regions are completely smoothed in the human atria. Consequently, the sheep exhibit significant micro-reentrant risk in the RAA and LAA, whereas the human atria show no risk in the LAA, and minimal structural detail at the entrance to the RAA. Conversely, where small structural holes have been sealed in the sheep atria, numerous small holes are found in the human atria, particularly in the superior right atrium and the PV sleeves, where extensive micro-reentrant risk is

observed. These regions are susceptible to micro-reentry even at high risk in the human where most of the OVF bias is lost in the sheep atria. This suggests that small inhomogeneities in the atrial structure have a significant impact on where micro-reentry is observed and how extensive the bias to regions with low OVF is.

6.6 Discussion

In this chapter, we have introduced a novel computational approach for analysing how the structural substrate for micro-anatomical reentry may develop as the atria accumulate diffuse interstitial fibrosis. The technique combines spatial networks with fibre tractography, and assumes that interstitial fibrosis accumulation can be modelled as the decoupling of a spatial network, a concept closely related to spatial percolation [192].

Many studies investigate the electro-anatomical basis for micro-anatomical reentry. Broadly speaking, such studies fall into three categories. In the first, the substrate can be identified directly, given sufficiently high resolution imaging, by observing the presence of a micro-anatomical driver. Examples include [144, 193], identifying several key factors in the formation of the reentrant substrate. Of particular note are the orientation of muscle fibres, the structure, thickness, and thickness gradients of the atria, and the accumulation of fibrosis, particularly interstitial fibrosis [212].

Attempting to quantify these factors, the authors in [193] applied optical mapping to explanted human atria to identify how specific values of wall thickness and fibrosis burden correlate to the location of the known micro-reentrant substrate [193]. Driver regions were found to correlate well with areas of 20-30% wall thickness and 20-30% fibrotic burden, notably the junction of the RIPV and PLA, and between the CT and fibrotically insulated PMs. Fibre misalignment was also implicated, specifically between the CT and the PMs which also exhibit abrupt changes in the local wall thickness. However, these studies took a largely static view of the micro-reentrant substrate, not considering how the substrate may change under electrical or structural remodelling.

In the second category, the substrate for micro-reentry has been established in the atria, but is hidden due to the lack of a clearly observable driver. One recent study has addressed this problem by noting that the visible micro-reentrant substrate varies strongly with variable atrial refractoriness [213]. Hence, the authors were able to demonstrate that the hidden micro-reentrant substrate could be unmasked and stabilised by shortening atrial refractoriness with adenosine. Such changes, induced pharmacologically, are analogous to some of the electro-mechanical changes that may be expected over time from atrial remodelling [141]. However, a robust framework for predicting the future micro-reentrant substrate is still lacking.

No study is yet to address the third category which is to predict how the substrate for micro-reentry will develop in the future on a patient specific basis, a problem directly related to the increasing need for arrhythmic risk stratification [214].

In this proof of concept – which we acknowledge is far from clinical applicability – we have attempted to take a step towards addressing this and ask how the substrate for micro-reentry may develop over time. Specifically, we ask how different parts of the atria are susceptible to micro-reentry at different characteristic fibrosis ranges; some regions of the atria, particularly those where the atrial walls are thin and there is significant fibre misalignment, may be susceptible to micro-reentry at a low fibrotic density, whereas other regions may require higher fibrosis densities before micro-reentry can be induced. Such an approach is currently only possible in-silico, primarily due to ongoing challenges with accurate in-vivo atrial imaging [215].

Our results implicate the role of wall thickness and the misalignment of fibres, suggesting that thin atrial walls and reduced longitudinal fibre coupling both enhance the probability that a region is susceptible to micro-reentry at lower fibrosis densities, supporting the findings in [193]. Moreover, our study suggests that the dependence on these factors evolves as the density of interstitial fibrosis grows in the atria. In particular, our study suggests that the spatial spread of the micro-reentrant substrate increases dramatically with small reductions in the spatial network coupling, and indicates that the bias to thin atrial regions with complex fibre morphology is reduced as the micro-reentrant substrate becomes more spatially diverse. Many

of the specific regions highlighted as risk substrates for micro-reentry in [193] and [213], such as PM ridges and the superior right atrium, naturally emerge as risk substrates using our method, most likely due to their position in thin atrial regions.

One explanation for the importance of wall thickness, fibrosis density, and local wall curvature is that driver regions emerge from percolation-like dynamics where the critical fibrosis density is dependent on the thickness of the structure studied [169, 216]. This in turn may result in micro-anatomical circuits anchoring to the atrial surfaces in paroxysmal AF but distributing across the atrial wall in persistent AF, as suggested in chapter 5. Aside from micro-reentry, modelling fibrosis using percolation-style distributions is known to perform particularly well when simulating patterns of AF maintenance [217], and has been used to explain complex fractionated electrograms and reentries in fibrotic border zones [168, 210].

In the wider AF literature, a number of computational studies consider the role of atrial structure on AF dynamics. In most cases, these studies investigate how structural factors affect electrical wavefront dynamics by solving the mono- or bidomain equation coupled with a suitable atrial cell model [216, 218–223]. However, even the most advanced patient-specific modelling methodologies suffer from the continued struggle to extract precise fibre orientation data and high resolution fibrosis profiles (particularly diffuse interstitial fibrosis) in-vivo [173].

Until these issues are resolved, computational modelling must continue to be used as a platform for hypothesis testing, studying the dynamics of AF initiation and maintenance with a variety of methods and across scales. In particular, the clearest insights may be found from models which contrast the role of a specific electro-anatomical feature with null models in the absence of that feature. Examples include models with and without realistic tissue geometries [218], with and without patient-specific fibrosis [224], isotropic vs. anisotropic fibre structures [221], and continuous vs. discrete modelling modalities [172]. We hope that the techniques introduced in this chapter can add to the range of techniques employed for hypothesis testing in cardiac electrophysiology.

6.6.1 Limitations

The aims of this study are highly focused, discussing the structural basis for micro-anatomical reentry only. Many other factors may affect the probability that a micro-reentrant substrate forms in a given region of the atria, including other forms of fibrosis and ion channel remodelling. Our approach does not consider these factors and cannot exclude their importance to micro-anatomical reentry, although from an ontological point of view, we believe it is important that the spatial network method is fully understood in its simplest form before adding complexities.

The key technical limitations of our work come under two main categories: (1) Imaging related limitations and (2) limitations arising from specific modelling choices in the construction of the spatial networks. In our view, the most important limitation relates to the datasets used for analysis. As noted throughout, the human fibre dataset has undergone extensive smoothing, lacking realistic local heterogeneity in the fibre structure; the observed heterogeneous regions are likely artefacts from the synthetic fibre generation method. Similarly, the atrial geometry is heavily smoothed, lacking fine structural detail in regions such as the RAA and LAA, but with numerous structural holes across the geometry. Given the strong evidence that these features influence the identified risk substrate, any future work must ensure that predictions are not simply an artefact of the data acquisition and generation process.

Although using high resolution fibre maps like those in [197, 198] may avoid these issues, acquiring such data is not feasible in-vivo, remaining a major ongoing research challenge [225, 226]. It is an open question whether sufficiently high quality fibre data can be acquired in-vivo for use with our method. Although simple methods such as those in [195] do not appear suitable for our method, other approaches for extracting estimates of fibre orientation heterogeneity in-vivo may have more potential [227].

The datasets used also lack information on individual-specific fibrosis distributions, knowledge of which would be necessary in a clinical setting. However, given the limitations of current LGE-MRI techniques for detecting precise fibrosis densities [215, 228], combining knowledge of the spatial distribution of fibrosis with our approach which focuses on specific fibrosis densities

will remain challenging. Future work should consider how best to combine these approaches.

Finally, our approach requires further validation before predictions can have real clinical relevance. We have demonstrated that the spatial networks accurately preserve the underlying properties of each atrial fibre map, and that our techniques ensure an even density of nodes across the network. We have shown that the simple discrete diffusion model (DDM) used to identify micro-reentry produces qualitatively realistic activation times in the isotropic and anisotropic cases, and that microscopic anisotropy values are consistent with experimentally measured values [229, 230]. However, for clinical relevance, it is important to validate the approach directly against experimental data, ideally with datasets with a history of micro-anatomical reentry. If possible raw data should be available such that fibre maps can be generated using a range of methods at different levels of smoothing, and that the influence of small structural holes can be accurately tested. If possible, data may be acquired at different levels of interstitial fibrotic density to assess the validity of our method's longitudinal predictions. If such an approach is not feasible, the method may be tested against advanced simulations of AF dynamics using the mono- or bidomain models in the presence of structural heterogeneities at high resolution.

6.7 Conclusions

In the previous two chapters, we have discussed two simple percolation-based models for simulating the formation of micro-anatomical reentrant circuits in the atria. Following a series of experimental studies, there is growing evidence that these circuits may be a plausible mechanism for initiating and maintaining atrial fibrillation [144, 193, 213].

Our approach to modelling follows a simple "less-is-more" philosophy. Rather than simulate precise atrial ion-channel dynamics in great detail, we implement electrical wavefront propagation on a 3d lattice in chapter 5, and a spatial network in chapter 6, using simple cellular automata style dynamics. This choice naturally limits our ability to study many of the rate-dependent effects which alter AF dynamics. However, the simplicity of our method places a

clear focus on the structural aspects of AF initiation and maintenance, and highlights how the accumulation of electrically insulating interstitial fibrosis (simulated by removing bonds in our lattice or spatial network) results in a changing spectrum of AF phenotypes driven by micro-reentry.

Although our physics based approach is still far from direct clinical relevance, the trajectory shown in the transition from the original 2d model developed in 2015 [145], through the 3d version discussed in chapter 5, and ending with the spatial network approach in chapter 6, demonstrates the potential of these methods for the future, and the value in simple modelling studies as the foundation for more advanced techniques.

In its current form, our approach corroborates a number of experimental findings and offers several novel insights which warrant further investigation. In particular, we demonstrate that the percolation model reproduces the clinically observed spectrum of AF as fibrosis accumulates, starting from unaffected healthy patients, to patients with short intermittent episodes (paroxysmal AF), and ending with persistent AF where fibrillation is maintained permanently.

As we move through this spectrum of AF risk, the models highlight the critical importance of tissue thickness and fibre structure, showing that micro-anatomical reentrant circuits are likely to form in regions where the atrial walls are thin and exhibit substantial fibre misalignment. We suggest that the dependence on these factors is complex, continuously evolving with the absolute level of micro-reentrant risk, demonstrating that, as fibrosis accumulates in the atria, there is a rapid increase in the absolute number of regions where a circuit can form. The spatial distribution of these risk regions becomes sparser and more evenly distributed across the atria as the overall risk level increases.

In my view, this is a critical result: for patients with paroxysmal AF, catheter ablation offers reasonable treatment success rates. Our results suggest that this is unsurprising – if AF episodes are intermittent, it is likely that only a very small number of regions are responsible for initiating and maintaining AF, and our results suggest these regions will be highly clustered and close to the endocardial surface. Such a small number of circuits should be relatively easy to isolate via ablation, preventing future AF episodes. However, if AF episodes are persistent and long-

lasting, it is far more likely that a large number of spatially diverse regions are responsible for initiating and maintaining AF. Even if all these regions could be isolated or destroyed by ablation, it would likely come at the cost of serious damage to the atrial muscle from the ablation lesions, and treatment success would only be maintained until further fibrotic deposit rendered new atrial regions susceptible to micro-reentry. This suggests that there is likely an upper limit to how effective catheter ablation can be for treating atrial fibrillation – taking the liberty to speculate, it is my view that we are already very close to this limit.

This does not mean that there is not value in developing patient-specific approaches to the treatment of AF. Considering the current work, if the spatial network approach can be developed to use other sources of patient-derived fibre orientation, acquired in-vivo, the approach may become useful for risk stratifying patients based on longitudinal predictions of their micro-reentrant substrates, giving an indication of which patients will, and will not, benefit from treatment. Such applications will certainly require improved in-vivo imaging of the human atria, but in the short to medium term, the framework may be useful for further hypothesis testing, and may be applied with synthetic atrial fibre orientation data as these techniques become more advanced.

6.8 Chapter summary

- We extend the 3d percolation model from chapter 5 to a spatial network framework. The spatial network approach is adaptable to real atrial geometries.

- Using three anatomically derived atrial datasets, we study the frequency and distribution of micro-anatomical reentrant circuits. This is modelled by slowly decoupling the spatial network, simulating the accumulation of diffuse interstitial fibrosis.

- When the rate of micro-reentry is low, we find that micro-anatomical reentrant circuits are few in number and highly clustered in thin atrial regions with reduced transverse coupling.

- As the rate of micro-reentry increases, the degree of clustering reduces, and micro-anatomical reentrant circuits are found across a wider fraction of the atria. The dependence to thin atrial regions and regions with reduced longitudinal coupling is reduced.

- I suggest that the spatial diversity of micro-anatomical reentrant circuits when the rate of reentry is high suggests that catheter ablation is unlikely to be successful in more serious cases of AF.

6.9 Technical details

6.9.1 Fibre tractography

Each atrial dataset provides a 3d vector field representing the best estimate for the local fibre orientation at each point in the 3d image of the atria. However, there is no information in the datasets regarding fibre connections. To apply our spatial network methods, we generate fibre tracts across the atrial structure from the underlying local fibre orientations. For a summary of standard fibre tractography approaches, principally pioneered in neuroscience, see [231].

In developing our tractography procedure, we experimented with two existing algorithms: Fibre Assignment by Continuous Tracking (FACT) and Evenly Spaced Streamlines (ESS). FACT [232] is often used for modelling purposes, whereas the ESS algorithm was developed in neuroscience for visualisation purposes. The principal differences between the methods are described in Table (6.4).

The principal limitation of each method is that in FACT, arbitrary fluctuations in the density of fibres exist (shown in Fig. 6.12). In ESS, the mean fibre length is far shorter than that in FACT. This motivated the use of an adapted version of ESS that prioritised the generation of long fibres over short fibres.

	ESS	FACT
Tracking Method	Step along local orientation with fixed step size d_{step}	Follow local orientation until a voxel boundary is reached
Density Regulation	Generation of a fibre terminated if closer than d_{sep} to a different fibre	Global cost function is applied after fibre tracking to selectively remove fibres
Limitation	Fibres generated were short	Arbitrary density fluctuations existed in the fibre map

Table 6.4: The key differences between the fibre tractography approaches FACT and ESS. Generation of a fibre in ESS had an extra termination condition compared to FACT.

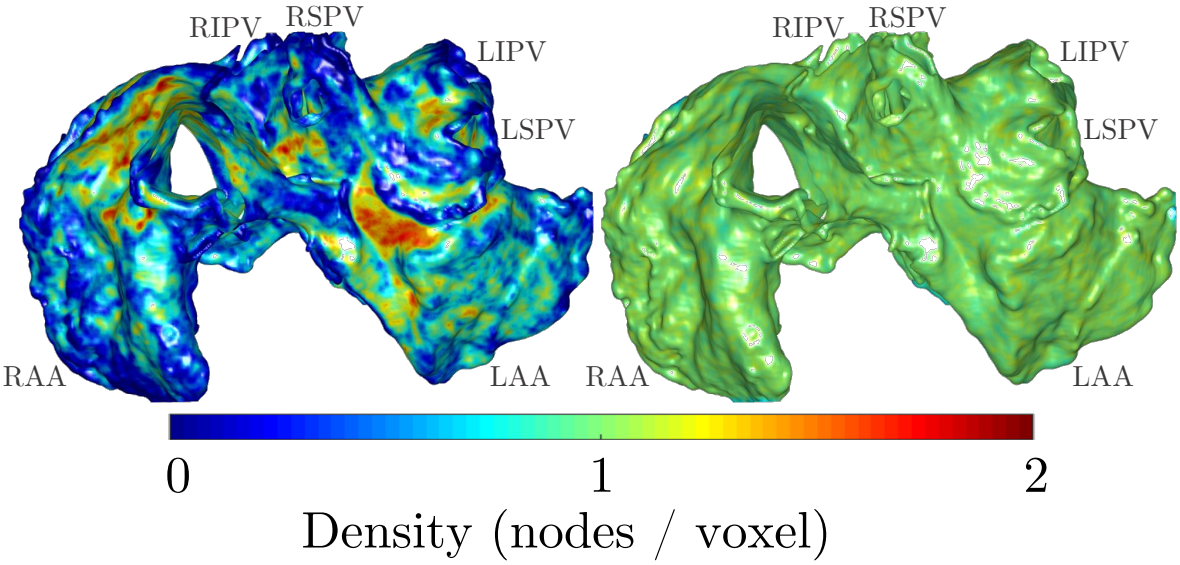

Figure 6.12: Heatmaps of node density, from networks generated on the healthy sheep atria using FACT (left) and ESS (right), shown from the superior view. The density regulation following FACT proved mostly ineffective, creating arbitrary regions of under/overdensity. ESS, by comparison, produced a network of near uniform density.

Given that no density information is supplied in the atrial datasets, we opted for an algorithm that favoured a uniform fibre density. By modifying the ESS algorithm, the fibre maps generated were of approximately uniform density and fibre length was maximised. This modified approach is as follows:

1. Seed points (an equal number to the number of occupied voxels) are generated randomly across the atrial structure in voxels not containing any existing fibre vertices.

2. Fibres are generated from seed points by Euler stepping along the local voxel orientation with fixed step size $d_{step} = 1$ in both the forwards and backwards directions.

3. For each fibre being propagated in the forward and backward direction, the fibre is terminated if, (a) the edge of the dataset is reached, (b) an empty voxel is reached, (c) the difference in angle between the current and next voxel orientation exceeds a threshold of 45°, see [233], (d) the distance between the new vertex and any vertex on a different fibre is less than $d_{sep} = 0.7$, see [201].

4. Once a fibre has terminated, it is added to the final fibre map if its length exceeds a threshold l, otherwise the fibre is discarded.

5. If all seed points have been checked and no new fibres exceeding the threshold length l can be added to the fibre map, steps 1-4 are repeated with a new threshold. In order, the thresholds used in each cycle are $l \in [30, 25, 20, 15, 10, 5, 0]$. This is a novel approach to fibre generation and aims to maximise atrial coverage while also ensuring that the fibres generated are as long as possible. The l values were chosen by trial and error, with the maximum length corresponding approximately to the longest fibres we observe using traditional tractography methods such as FACT.

6. Once a threshold of $l = 0$ has been reached, steps 1-4 are repeated until vertices placed on existing fibres are present in at least 99% of the voxels occupied by tissue in each dataset.

Some applications in computational cardiology, which implement a fibre structure in the atria for modelling purposes, explicitly build in specific known fibre bundles which are important for electrical propagation in the atria [234]. For our application, we do not give any special treatment to specific fibre bundles, relying solely on the local fibre orientation data in each dataset.

6.9.2 Discrete diffusion model

To identify the locations in the spatial network capable of forming micro-anatomical reentrant circuits, we search the network structure for an isolated fibre tract, defined as a continuous string of transversely disconnected nodes, such that a non-backtracking path from the start of the fibre tract, through the fibre tract, and back around the fibre tract exists and is at least length $\tilde{\tau}$, where $\tilde{\tau}$ is the refractory wavelength.

Computationally, such locations can easily be identified by applying a discrete diffusion model (DDM) to the network. The approach is based on previous simple models of micro-anatomical reentry in [3, 145]. For simulation purposes we set $\tilde{\tau} = 83$ such that the threshold length for a micro-reentrant circuit is marginally smaller than the circuits identified in [144] [3]. Note, however, that this does not exclude circuits forming that are significantly larger than the threshold, $\ell \gg \tilde{\tau}$, only preventing excessively small circuits of length $\ell < \tilde{\tau}$. The specific value of $\tilde{\tau}$ is largely irrelevant since increases (decreases) in $\tilde{\tau}$ can be counteracted by an appropriate decrease (increase) in the characteristic coupling c [8].

First, we define a point in the right atrium as the sino-atrial node, around which all nodes within an arbitrary radius 2 are automatically excited every $T = 1000$ time steps. All nodes which are in the excited state at time t_i enter the refractory state at $t_i + 1$, during which time nodes cannot be re-excited, and remain in this state for τ time steps. At time $t_i + \tau + 1$, these nodes enter the resting state. Nodes in the resting state at time t_i enter the excited state if (1) at $t_i + 1$ they are directly connected to a node which is excited at time t_i, and (2), conduction is not blocked. Nodes which should enter the excited state at time $t_i + 1$ may fail to do so with small probability $\epsilon = 0.0025$. This simulates potential unidirectional conduction block at the micro-level in a simple and efficient way, while avoiding extensive wavefront breakup (i.e., wavefronts should remain approximately planar as is expected during sinus rhythm). The specific value of ϵ is arbitrary, but must be sufficiently small such that in any given activation cycle, the probability of a micro-anatomical reentrant circuit forming is small [8].

[3] Starting with voxels with $100\mu m$ spacing, a value of $\tilde{\tau} = 250$ gives a minimum refractory wavelength of 25mm. To ensure a fixed refractory wavelength, $\tilde{\tau}$ must be divided by a factor of 3, $\tilde{\tau} = 250/3 \approx 83$, when the vector field is down-resolved to $300\mu m$ resolution.

A micro-reentrant circuit is identified by allowing the excited wavefront to diffuse through the spatial network from the sinus node, and observing where isolated stochastic conduction blocks result in the formation of a circuit, shown schematically in Fig. 6.13. This emerges naturally in the DDM, and is identified by locating the first node which activates multiple times within a single sinus activation cycle of length T. We have visually verified that the location identified computationally corresponds to the location in which we observe rapidly firing micro-reentrant triggers in the simulations.

As a result of the stochastic nature of conduction block, a potential reentrant circuit only has a small probability of triggering with each sinus wavefront which crosses it. Therefore, for simulation purposes we run the discrete diffusion model for 10^5 time steps to allow sufficient time for a reentrant circuit to activate if present. Once a circuit is identified, the simulation is terminated. In the case where multiple circuits coexist within a single spatial network, the first to activate is recorded, although this is largely mitigated by an appropriately small choice of ϵ. By repeating this process until at least 10^3 circuits have been identified for each fibre map and each value of the characteristic coupling distance c, we can construct a probabilistic risk map for the micro-reentrant substrate across the atria.

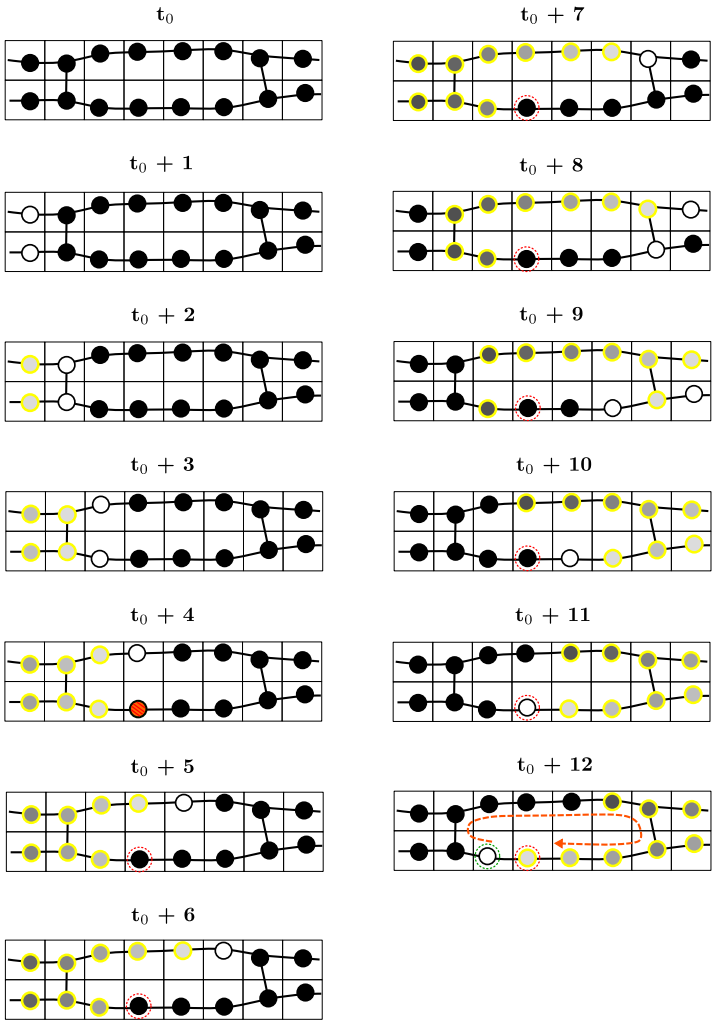

Figure 6.13: A schematic illustrating the mechanism by which a micro-reentrant substrate is detected by the discrete diffusion model. A dormant structure is present in the spatial network at time t_0 (black circles). A wavefront enters from left to right (white nodes). Excited nodes remain refractory for τ time steps after excitation (greyscale with yellow border). The wavefront is spontaneously blocked in the lower fibre at $t_0 + 4$ (orange-red checked circle). The wavefront proceeds along the top fibre and reenters the lower fibre at time $t_0 + 8$. If the closed loop is sufficiently long to form a micro-reentrant circuit, at time $t_0 + 12$, the node to the left of the previously blocked node will excite (circled in green), $\ell \geq \tilde{\tau}$. Otherwise, if the loop is too short, the node will be in the refractory state and will not excite, $\ell < \tilde{\tau}$. The DDM detects the repeated activation of the node circled in green as the location of a micro-reentrant substrate.

6.9.3 Atrial wall thickness & occupied voxel fraction

In this section we calculate the atrial wall thickness and the corresponding OVF values. The two distributions are not easy to compare since thickness is not defined in the bulk using our method. However, we illustrate that the OVF naturally increases in thicker regions of the atria, and that the two measures behave differently in regions of high wall curvature.

Calculated OVF and thickness

Figure 6.14 shows the atrial wall thickness values for the three atrial datasets calculated using the techniques developed in [235]. For comparative purposes, all three datasets are plotted using the same colourbar. Regions with a thickness in excess of 4mm are shown in grey to ensure high fidelity in the visualisation of thinner atrial regions.

Figure 6.15 shows the OVF calculations for the three atrial datasets. Visual inspection shows that thickness values fluctuate over a much larger range than the OVF calculations which appear smoother.

Figure 6.16 shows the OVF for an atrial cross section in each of the three atrial datasets. The figure demonstrates that the largest OVF values are found in the atrial bulk, and that the OVF naturally increases (decreases) with increasing (decreasing) wall thickness. For voxels further than 1.5 mm away from the atrial surfaces, the OVF is measured as 1. This demonstrates that the OVF captures the qualitative properties of traditional thickness measurements and extends naturally to the atrial bulk.

Apart from issues with defining thickness in the bulk, thickness measurements show small inconsistencies in regions of high curvature such as the pectinate muscles or vessel openings, see Fig. 6.17. These have low OVF but a large thickness, since the thickness measurement attempts to calculate the path from the endocardium to the epicardium. If using the thickness to quantify regions of micro-reentrant risk, spatially isolated fibres would counter-intuitively be identified as regions of high thickness. Using the OVF as an alternative measure avoids this issue.

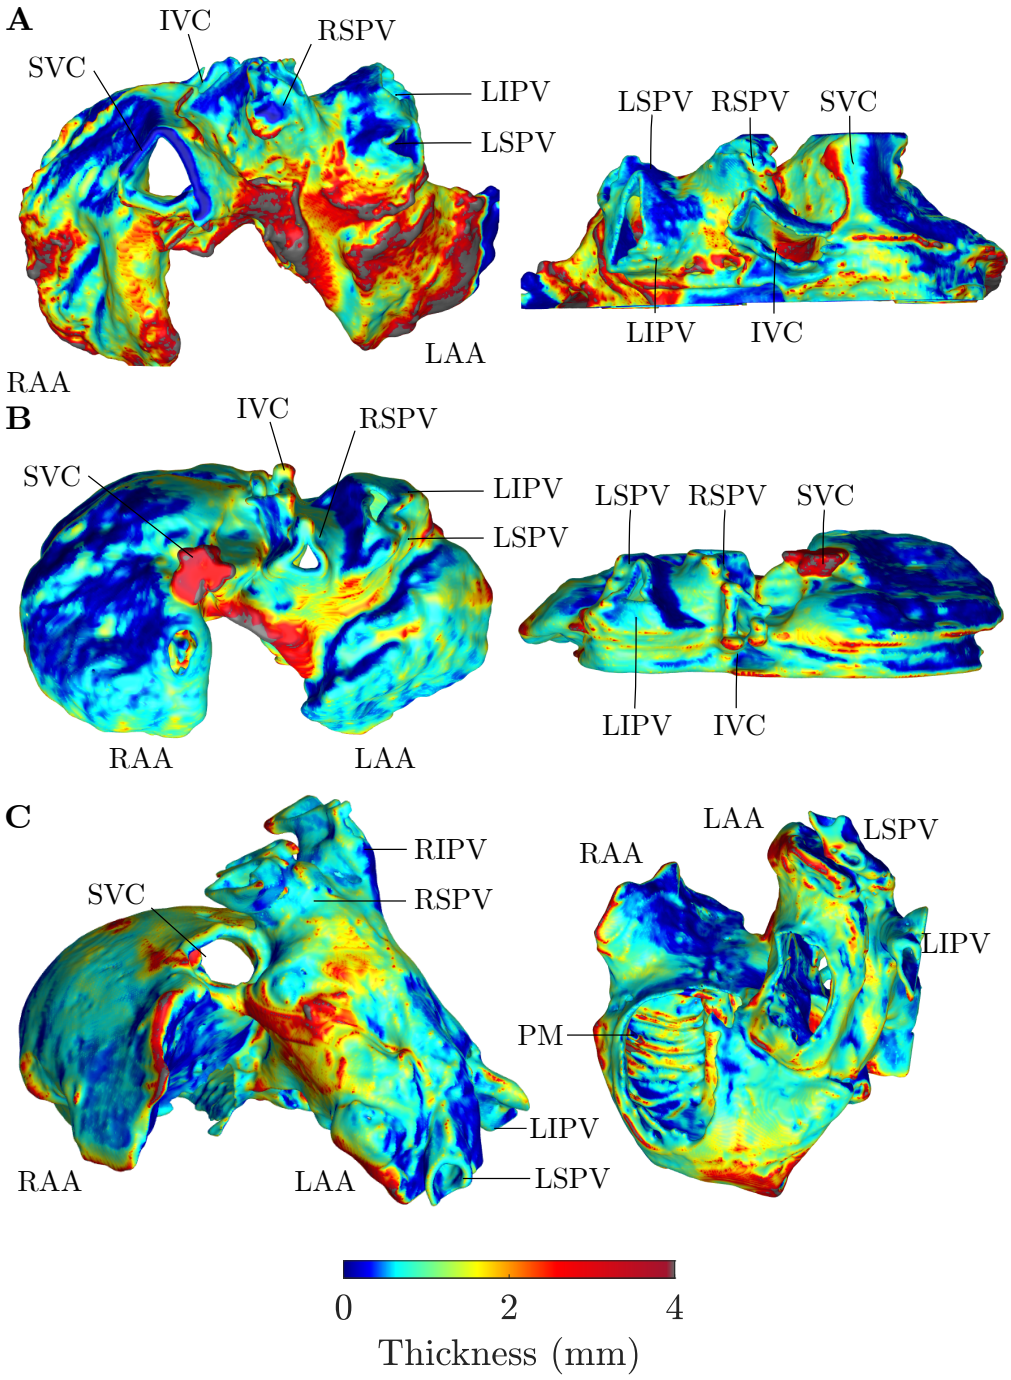

Figure 6.14: The thickness of the atrial wall for (a) the healthy sheep atria, (b) the heart failure sheep atria, and (c) the human atria. The colour scale is capped at 4mm with thicker regions shown in grey. Regions of very high curvature in some cases appear as points with very large wall thickness.

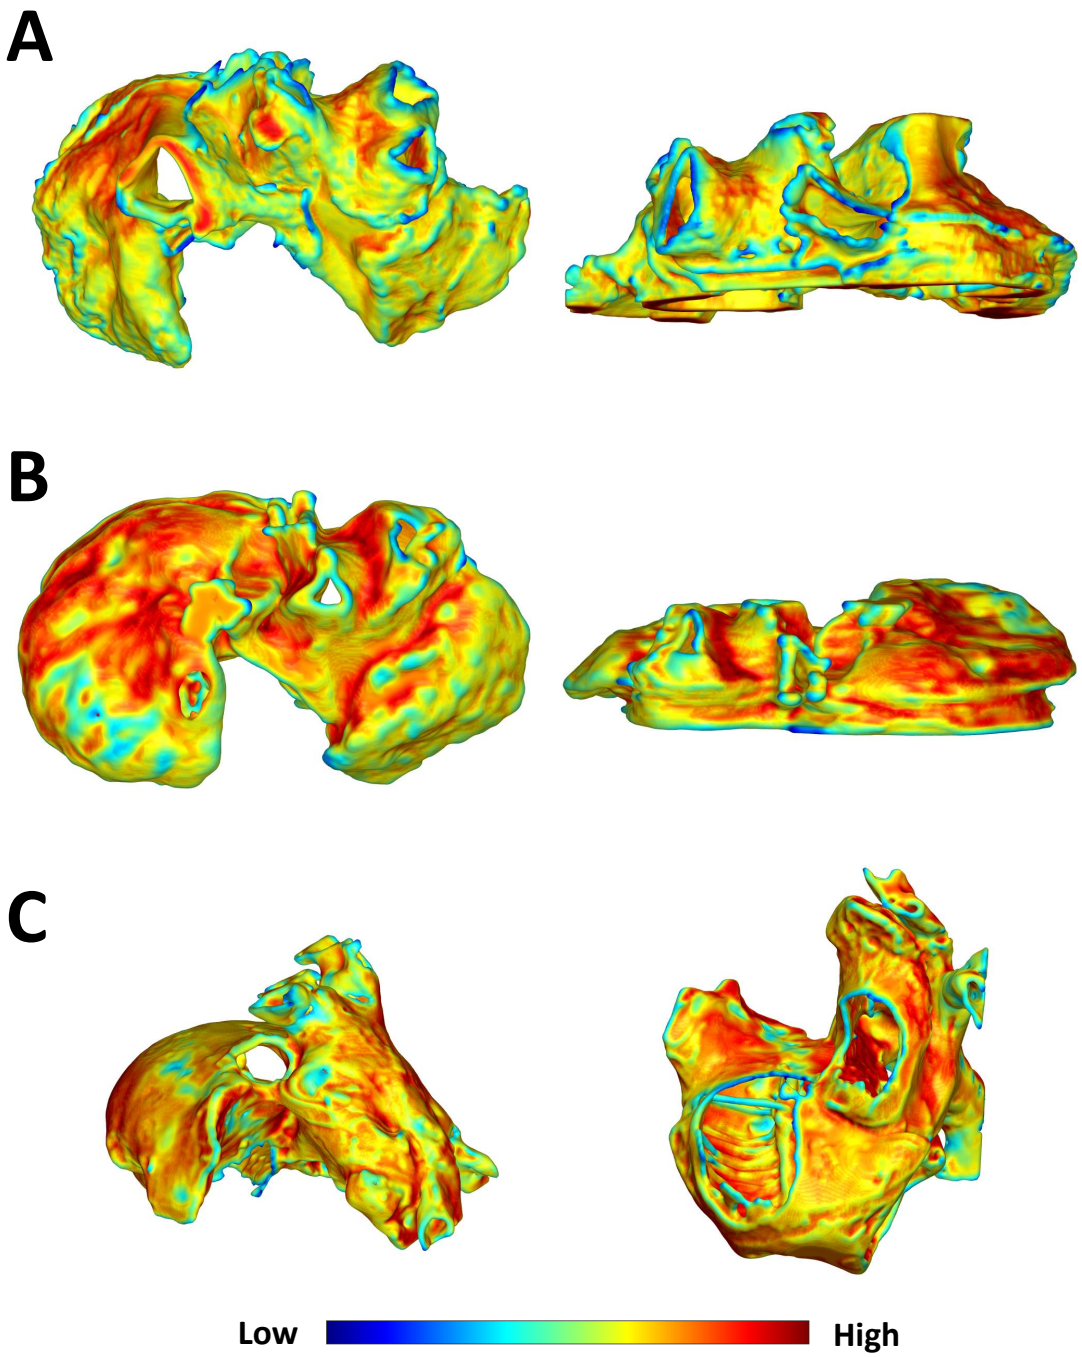

Figure 6.15: The OVF for (a) the healthy sheep atria, (b) the heart failure sheep atria, and (c) the human atria. Regions with the highest OVF values are not visible since they are within the atrial bulk.

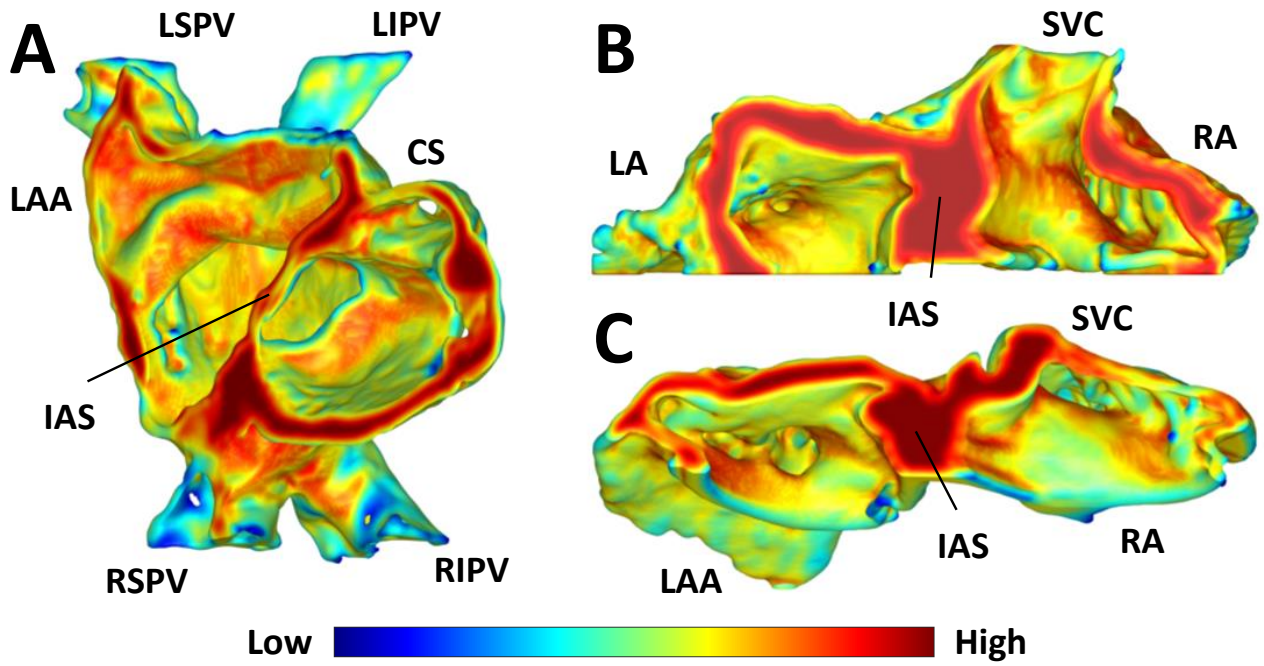

Figure 6.16: Three atrial cross sections showing the occupied voxel fraction (OVF) in the (A) human atria, (B) healthy sheep atria, and (C) the HF sheep atria. The figure shows that the largest OVF values are found in the atrial bulk, with lower values on the atrial surfaces. Bulk voxels more than 1.5 mm from the surfaces all have an OVF of 1. Regions of the atria which are thin (e.g. LAA in human atria) have a lower OVF than thick regions (e.g. IAS in sheep atria).

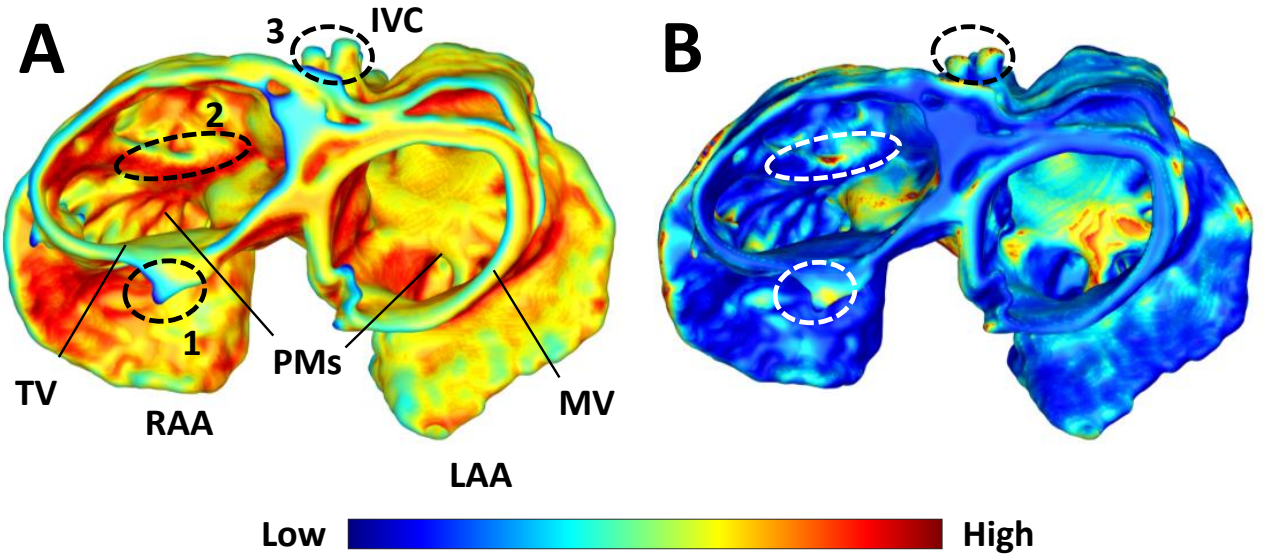

Figure 6.17: The (A) OVF and (B) wall thickness calculated in the heart failure sheep, shown from the inferior view. Regions which are spatially isolated and have significant wall curvature (regions 1–3), naturally appears as having low OVF. However, the wall thickness implementation results in these regions being measured as having a high thickness. This is particularly important along the PMs, at the TV and MV openings, and at the end of the IVC and PVs.

6.10 Supplementary results

6.10.1 Spatial distribution of risk substrate in the sheep atria

Figure 6.18 shows the spatial distribution of the micro-anatomical reentrant substrate for the healthy sheep atria at low, medium and high risk. The equivalent fibre-less null model is shown in Fig. 6.19, and additional views including atrial cross-sections are shown for both the fibre and fibre-less null models in Fig. 6.20. The corresponding figures for the HF sheep are shown in Fig. 6.21 (fibre model), Fig. 6.22 (fibre-less null model), and Fig. 6.23 (additional views).

Fibre model

At low risk, micro-reentrant circuits are almost exclusively clustered around the sleeve of the inferior vena cava (IVC). At medium risk, both sheep atria exhibit significant risk around the IVC and the RAA where the endocardium is dominated by ridged PMs. In contrast, the risk of reentry in the right and left atrial bulk, where the endocardium is smooth, is approximately zero. Particularly in the healthy sheep atria, a significant substrate arises along the PV sleeves, whereas for the heart failure sheep the risk along the PVs is non-zero, but not high. In the healthy sheep atria, there is noticeable risk at the opening of the CS, although this substrate is absent in the heart failure sheep atria, likely because the CS is sealed in this dataset.

Finally, at high risk the spatial distribution of the susceptible substrate is widespread for both sheep atria. Significant risk is observed along the IVC, at the PVs, and in both the right and left atria, including both atrial appendages where there are a large number of ridged PMs, although risk remains largely absent in regions where the endocardium is dominated by smooth tissue. Risk is maintained in the healthy atria at the opening of the CS, but is absent in the heart failure sheep.

Fibre-less null model

In agreement with the results for the human atrial dataset, the spatial distribution for the risk of micro-reentry is similar between the fibre and fibre-less null models for the sheep atrial datasets.

For the healthy sheep at low risk, the key risk region in both the fibre and fibre-less null models lies along the sleeve of the IVC. However, at medium and high risk, the fibre model shows diffuse risk across the atrial appendages whereas these risk regions are largely absent in the fibre-less null models. Conversely, the fibre-less null model shows extensive risk along the posterior right atrial wall and the opening of the SVC, whereas these regions are largely suppressed in the fibre model. Finally, the opening of the LIPV is identified as a key risk region in the fibre model, but risk in the fibre-less null model remains focused around the IVC at all risk levels, with only minimal risk observed at the LIPV.

For the HF sheep, the differences between the fibre and fibre-less null models are more pronounced, and there are noticeable differences between the healthy and HF sheep. At low risk, the IVC is the principal risk region in the fibre model, whereas the fibre-less null model shows extensive risk along an isolated tract of tissue bordering the TV. At medium and high risk, both the fibre and fibre-less models show risk across the atrial appendages. However, in the fibre model this risk is moderate and relatively spatially confined, whereas the fibre-less null model exhibits significantly higher risk across a wider region, particularly across the RAA. This is the opposite effect to the healthy sheep where diffuse risk is observed across the atrial appendages for the fibre model, but is largely absent in the fibre-less null model. The most likely explanation for these differences is the observation that the HF sheep is significantly thinner than the healthy sheep atria, see Fig. 6.14, although differences in the resolution at which the datasets are acquired may also play a factor.

Other differences between the fibre and fibre-less null models for the HF sheep include the extensive risk along the right endocardial wall in the fibre-less null model which is suppressed in the fibre model. The fibre-less null model also demonstrates more extensive risk across the

PV sleeves, whereas risk is low in these regions for the fibre model.

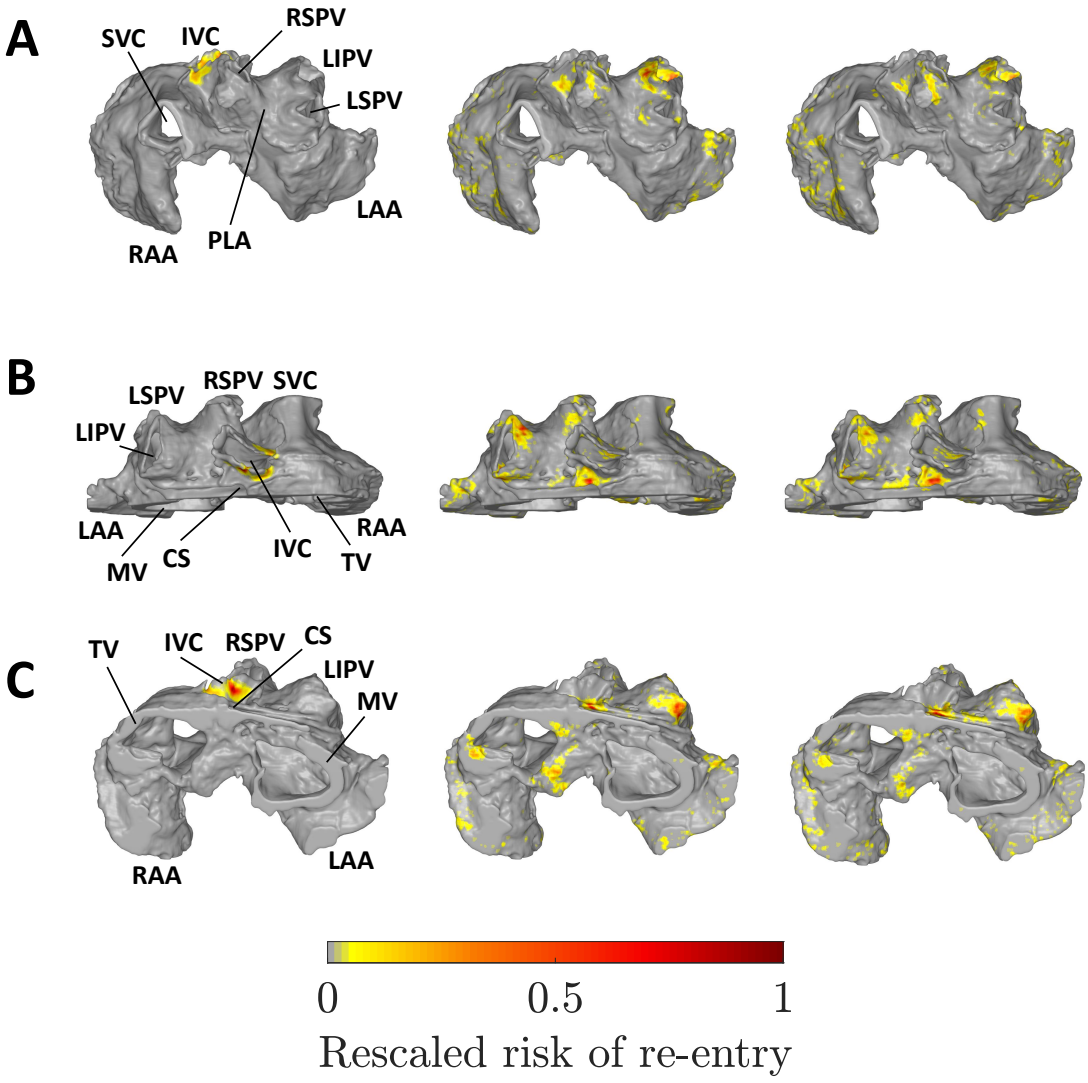

Figure 6.18: The spatial distribution of the micro-reentrant substrate in the fibre model at low (left), medium (middle) and high risk (right) for the healthy sheep atria. A: Superior view. B: Posterior view. C: Inferior view. IVC/SVC: Inferior/ superior vena cava. PLA: Posterior left atrium. RAA/LAA: Right/ left atrial appendage. TV/MV: Tricuspid/ mitral valve opening. CS: Coronary sinus. RIPV/LIPV/RSPV/LSPV: Right/ left, inferior/ superior, pulmonary vein.

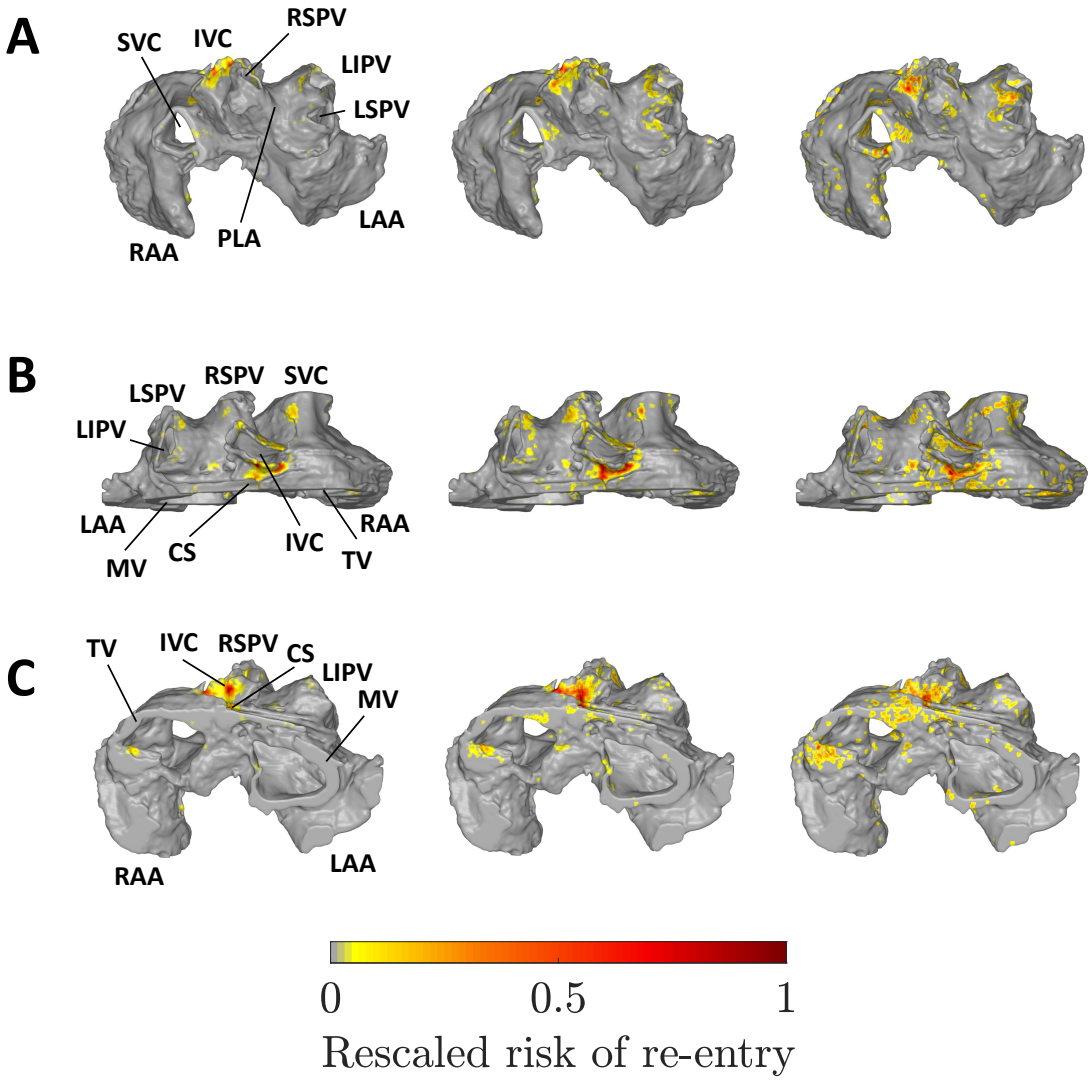

Figure 6.19: The spatial distribution of the micro-reentrant substrate in the fibre-less null model at low (left), medium (middle) and high risk (right) for the healthy sheep atria. A: Superior view. B: Posterior view. C: Inferior view. IVC/SVC: Inferior/ superior vena cava. PLA: Posterior left atrium. RAA/LAA: Right/ left atrial appendage. TV/MV: Tricuspid/ mitral valve opening. CS: Coronary sinus. RIPV/LIPV/RSPV/LSPV: Right/ left, inferior/ superior, pulmonary vein.

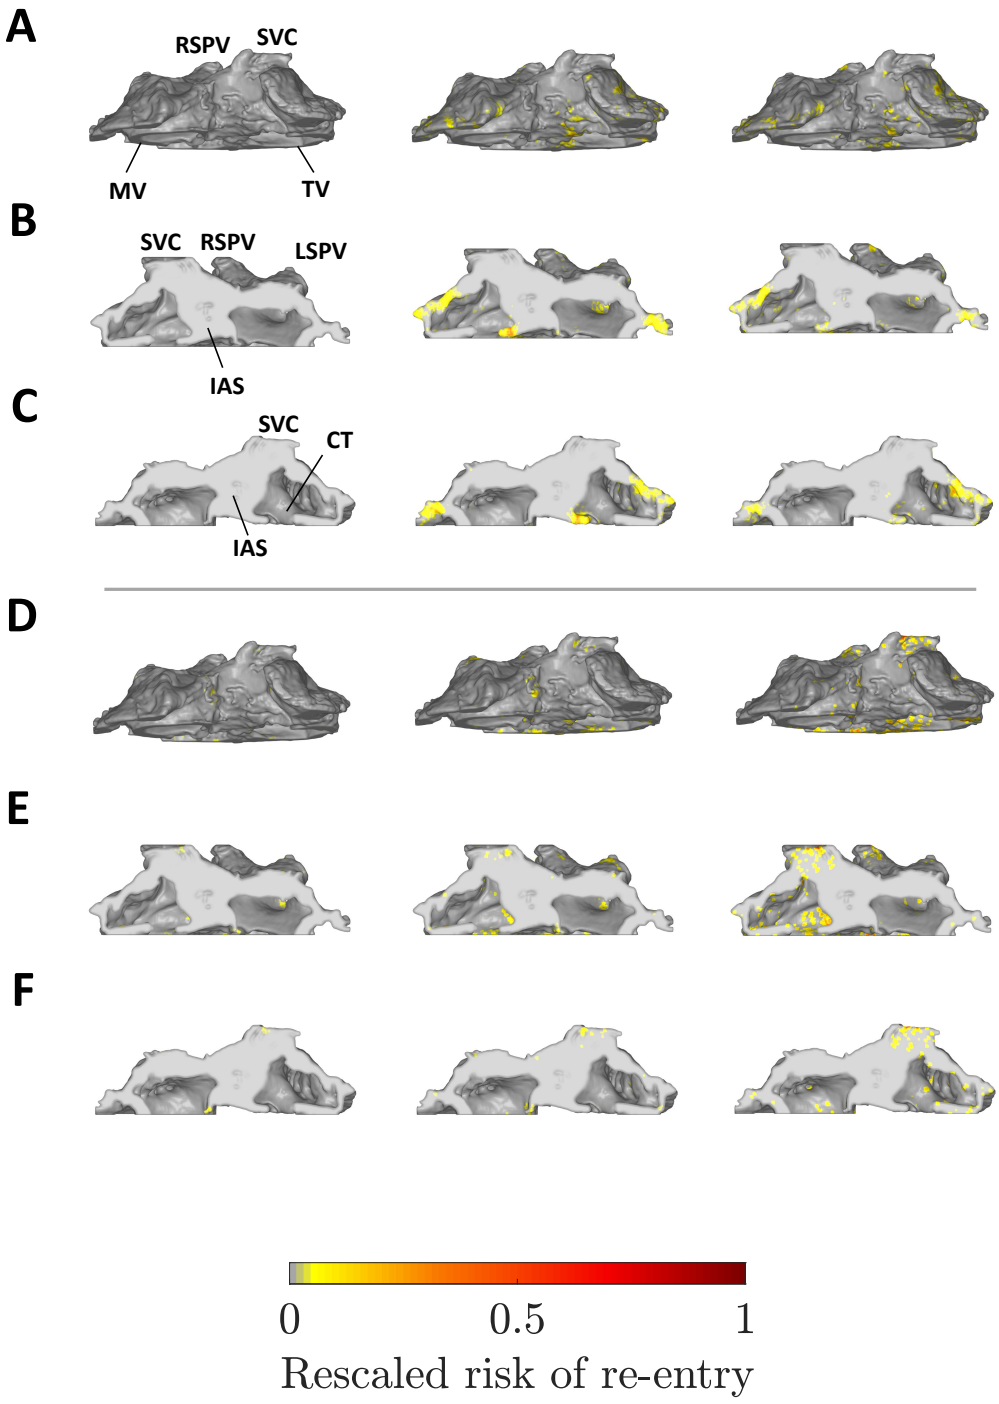

Figure 6.20: The spatial distribution of the micro-reentrant substrate for the healthy sheep atria in the fibre (A-C) and fibre-less null models (D-F) at low (left column), medium (middle) and high risk (right). A/D: Anterior view. B/E: Cut-through anterior view. C/F: Cut-through posterior view. SVC: Superior vena cava. CS: Coronary sinus. RSPV/LSPV: Right/ left, superior, pulmonary vein. IAS: Inter-atrial septum. CT: Crista terminalis. TV/MV: Tricuspid/ mitral valve opening.

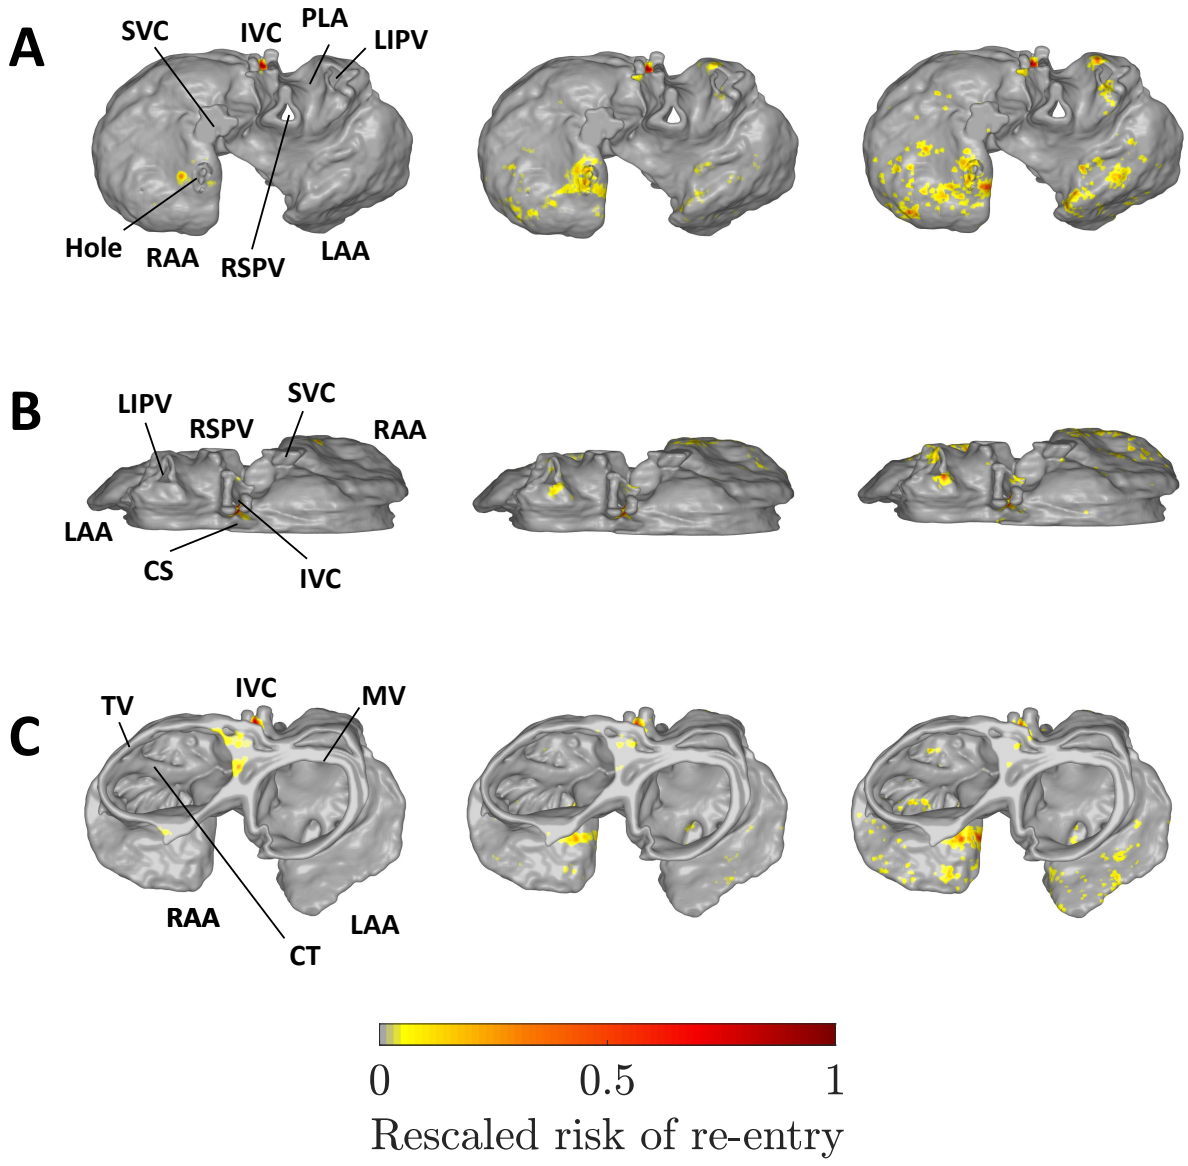

Figure 6.21: The spatial distribution of the micro-reentrant substrate in the fibre model at low (left), medium (middle) and high risk (right) for the HF sheep atria. A: Superior view. B: Posterior view. C: Inferior view. IVC/SVC: Inferior/ superior vena cava. PLA: Posterior left atrium. RAA/LAA: Right/ left atrial appendage. TV/MV: Tricuspid/ mitral valve opening. CS: Coronary sinus. RIPV/LIPV/RSPV/LSPV: Right/ left, inferior/ superior, pulmonary vein. The hole labelled in A is an artefact of imaging and is not an opening in the dataset.

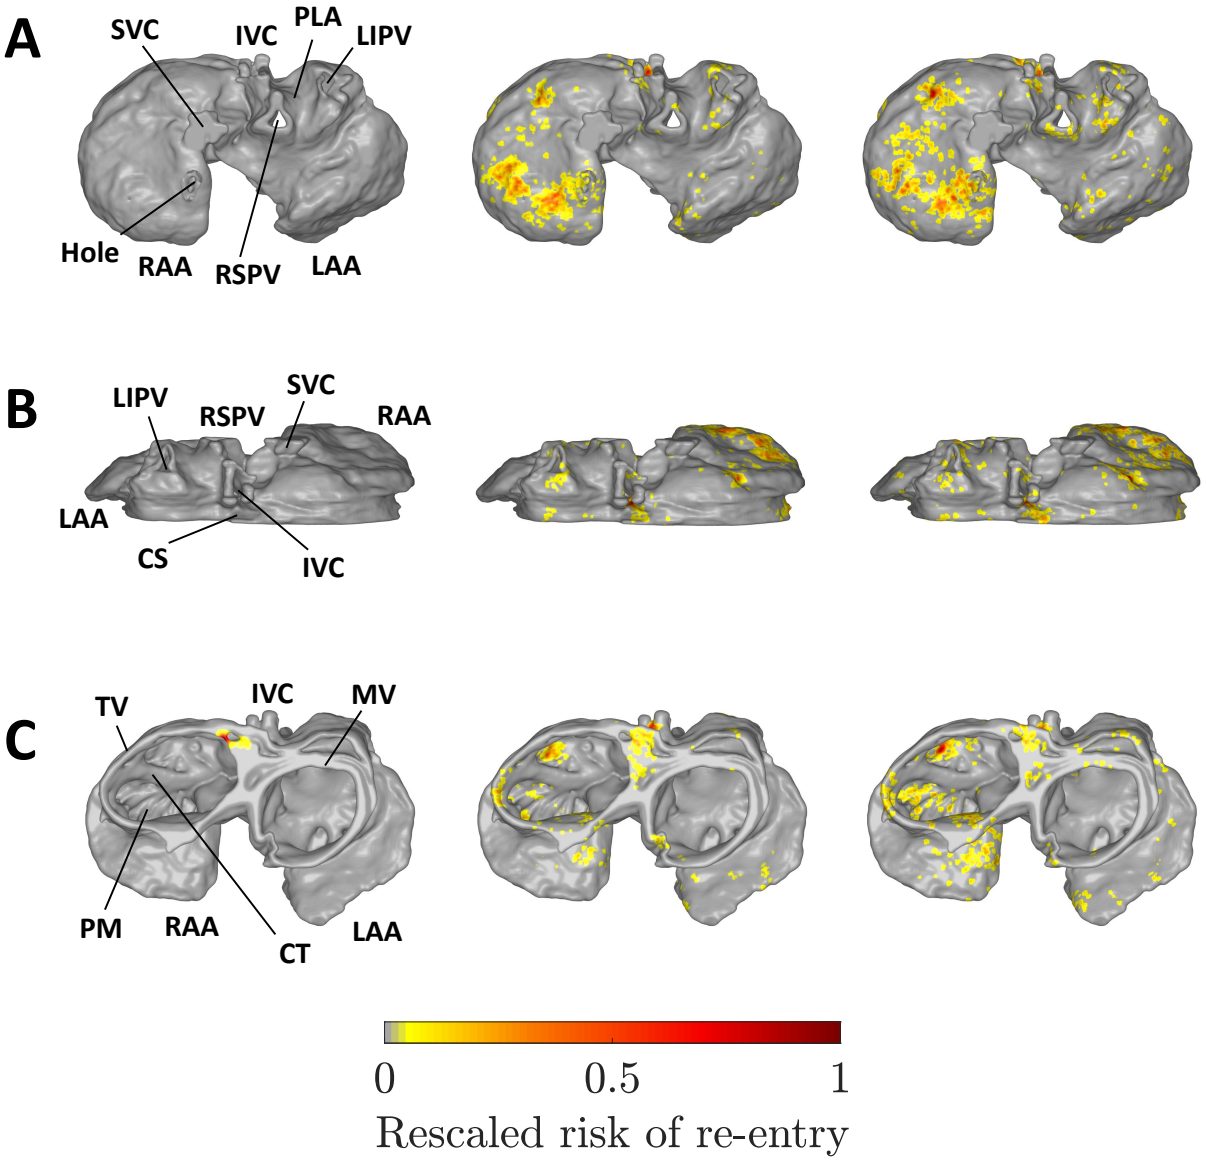

Figure 6.22: The spatial distribution of the micro-reentrant substrate in the fibre-less null model at low (left), medium (middle) and high risk (right) for the HF sheep atria. A: Superior view. B: Posterior view. C: Inferior view. IVC/SVC: Inferior/ superior vena cava. PLA: Posterior left atrium. RAA/LAA: Right/ left atrial appendage. TV/MV: Tricuspid/ mitral valve opening. CS: Coronary sinus. RIPV/LIPV/RSPV/LSPV: Right/ left, inferior/ superior, pulmonary vein.

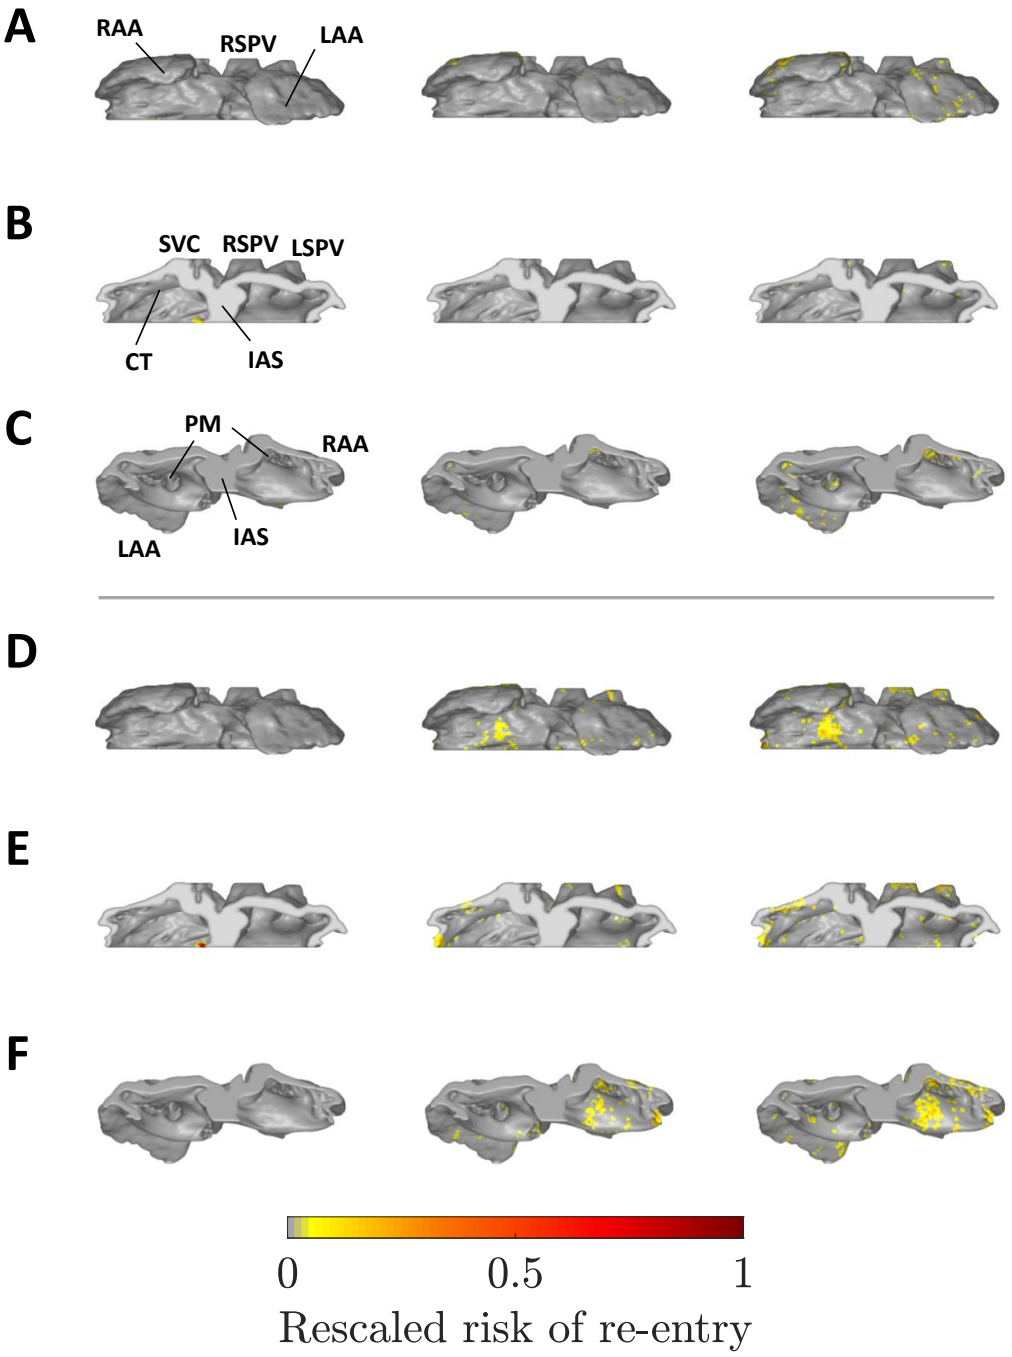

Figure 6.23: The spatial distribution of the micro-reentrant substrate for the heart failure sheep in the fibre (A-C) and fibre-less null models (D-F) at high (left column), medium (middle) and low coupling (right). A/D: Anterior view. B/E: Cut-through anterior view. C/F: Cut-through posterior view. SVC: Superior vena cava. CS: Coronary sinus. RSPV/LSPV: Right/ left, superior, pulmonary vein. IAS: Inter-atrial septum. CT: Crista terminalis. RAA/LAA: Right/ left atrial appendage. PM: Pectinate muscles.

6.10.2 Effect of fibre structure in sheep atria

Figures 6.24 and 6.25 show how the distribution of OVF and LCF values in the healthy and HF sheep atria respectively affect the distribution of micro-reentrant risk. Each figure shows the distribution where each voxel is weighted equally, weighted according to the risk of micro-anatomical reentry for the fibre model, and weighted according to the risk in the fibre-less null model, each at low and high overall risk. The figures are qualitatively consistent with the data for the human atria.

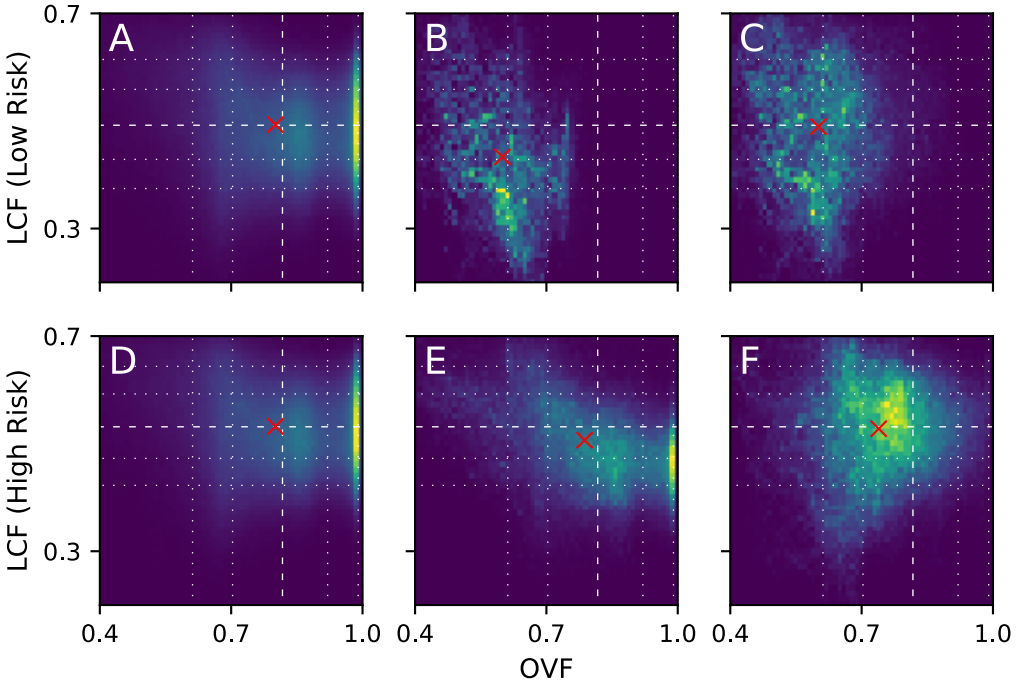

Figure 6.24: The distribution of LCF and OVF values in the healthy sheep atria at low (A-C) and high (D-F) risk. (A) & (D) The distribution across all the voxels in the healthy sheep atria. (B) & (E) The distribution weighted by the relative risk of reentry in each voxel for the fibre model. (C) & (F) The equivalent for the fibre-less null model. Note, the LCF is not defined for the fibre-less null model, but for illustrative purposes we include the LCF for the equivalent fibre model. Dashed lines: median. Dotted lines: 10^{th}, 25^{th}, 75^{th} and 90^{th} percentiles. Red crosses: weighted mean.

Note that in particular for the healthy sheep atrium, the densest region of the OVF and LCF distribution where voxels are uniformly weighted (Figs. 6.24(A) and (D)) is at OVF ~ 1. This indicates that the majority of voxels in the healthy sheep atria are far from the atrial walls. This is not the case for the human atria or the HF sheep atria.

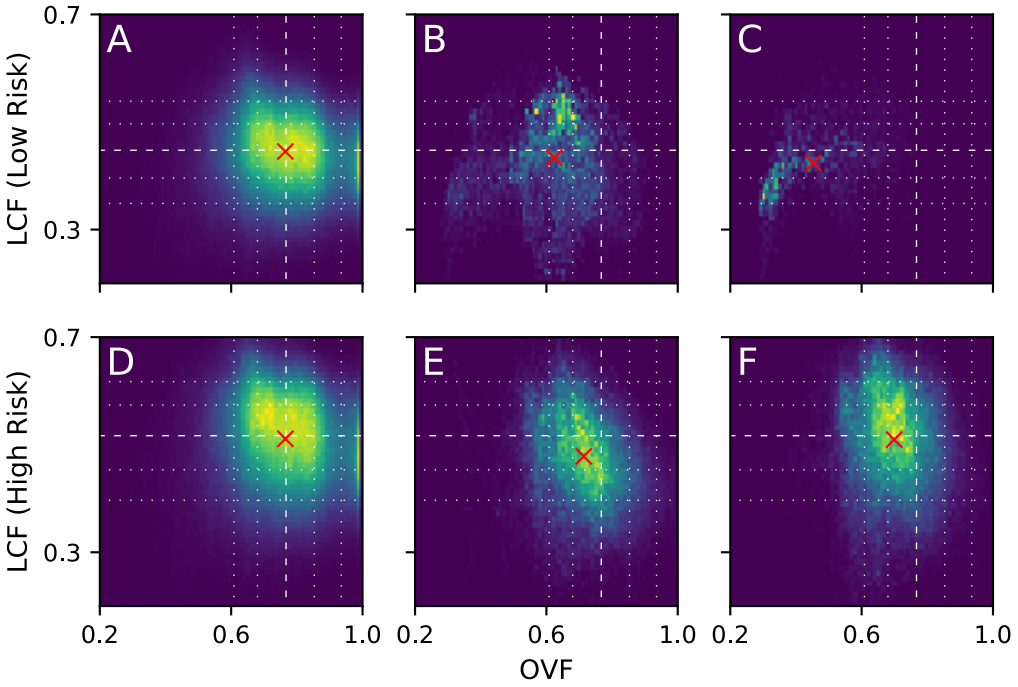

Figure 6.25: The distribution of LCF and OVF values in the HF sheep atria at low (A-C) and high (D-F) risk. (A) & (D) The distribution across all the voxels in the human atria. (B) & (E) The distribution weighted by the relative risk of reentry in each voxel for the fibre model. (C) & (F) The equivalent for the fibre-less null model. Note, the LCF is not defined for the fibre-less null model, but for illustrative purposes we include the LCF for the equivalent fibre model. Dashed lines: median. Dotted lines: 10^{th}, 25^{th}, 75^{th} and 90^{th} percentiles. Red crosses: weighted mean.

For the low risk cases in both sheep, the mean of the distributions weighted by both the fibre and fibre-less null models falls at significantly lower OVF values than the uniform case. In particular, no risk of micro-reentry is observed in regions where the OVF is above the 50th percentile, suggesting a complete suppression of micro-reentry in thick regions of the atria.

At high risk, the bias to low OVF values is significantly reduced, but a small bias is still visible. In the healthy sheep fibre model, risk is observed at an OVF ~ 1, but only in regions with reduced LCF. In contrast, micro-reentry is completely suppressed in regions with OVF ~ 1 in the fibre-less null model. This indicates that reduced longitudinal fibre coupling enables micro-reentry in thicker regions of the atria at high risk.

6.10.3 Statistical analysis of OVF and LCF distributions

To test for statistically significant differences between the LCF and OVF distributions presented above, we apply a permutation test [211]. Permutation tests are non-parametric and therefore avoid issues that arise due to the non-normality of the OVF distributions for each dataset. LCF values are mostly normally distributed, but we use the permutation test in both cases for consistency.

For Gaussian distributions, a standard measure for the effect size quantifying the difference between two distributions is Cohen's d, defined as

$$d = \frac{\bar{X} - \bar{Y}}{s}, \tag{6.3}$$

where $\bar{X}$ and $\bar{Y}$ are the mean values of the two distributions, and s is the pooled standard deviation. Since the OVF distributions are non-normal, we apply the nonparametric analogue of Cohen's d by using the *median absolute deviation*, defined as

$$MAD(X) = C \cdot med\left(|X_i - med(X)|\right), \tag{6.4}$$

where X_i is the i-th sample from the distribution X, $med(\cdot)$ denotes the median, and $C = 1.4826$ is a consistency constant that ensures the nonparametric Cohen's d maps to the regular Cohen's d for normal distributions. Using this form, we can also define the *pooled median absolute deviation* as

$$PMAD(X, Y) = \sqrt{\frac{(n_X - 1) \cdot MAD^2(X) + (n_Y - 1) \cdot MAD^2(Y)}{n_x + n_Y - 2}}. \tag{6.5}$$

Hence, we reach the non-parametric equivalent of Cohen's d by replacing the means in Eq. (6.3) with the medians, and replacing the pooled standard deviation with $PMAD(X, Y)$, written as

$$d_{NP} = \frac{med(X) - med(Y)}{PMAD(X, Y)}, \tag{6.6}$$

which reduces to Cohen's d for a normal distribution. For more details on this method see

[236].

Statistical analysis of OVF and LCF distributions

Tables 6.5, 6.6, and 6.7 summarise the statistical impact of the occupied voxel fraction in the human, healthy sheep and HF sheep respectively. All tests are carried out using a permutation test with 10,000 random permutations of each pooled distribution. The null hypothesis is that the two tested samples are drawn from the same underlying distribution. The equivalent tests are shown for the LCF distributions in Tables 6.8, 6.9 and 6.10

In each table we compare the OVF (or LCF) distribution for voxels identified as the micro-reentrant substrate in the fibre model at low, medium and high risk, against an equally sized sample of voxels drawn uniformly at random from throughout the atrial geometry. The equivalent is provided for each of the fibre-less null models, tested against uniform sampling. Finally, we test the fibre models against their corresponding fibre-less null models at low, medium and high risk.

For the OVF, we find statistically significant differences ($p < 0.01$) for all distribution comparisons, with the exception of the high risk fibre model compared to the uniform distribution in the healthy sheep atria (see Table 6.9) where we find $p = 0.06$.

Focusing on the comparison between the risk distributions and the uniform distributions, we note that the effect sizes are largest for the low risk cases, moderate for the medium risk cases, and lowest for the high risk cases, demonstrating the reduction in the bias to low OVF values with increasing risk. Despite this, even at the highest risk levels a small bias to reduced OVF values is retained (except for high risk fibre model in healthy sheep). Comparing the fibre models to the fibre-less null models, at medium and high risk, there is evidence that the bias to small OVF values is stronger in the fibre-less null models than the fibre models.

For the LCF, we find statistically significant differences for all distribution comparisons with the exception of: (1) human null high risk vs. uniform, healthy sheep null low risk vs. uniform, HF sheep null medium risk vs. uniform, and HF sheep null high risk vs. uniform. We do not

expect any of the null distributions to be significantly different from the uniform case since the LCF should have no effect in the fibre-less null models. Any differences which are observed are indirect and are most likely due to small correlations between OVF and LCF values. In particular, for spatially isolated ridges such as the PMs, we often find that these regions of low OVF also correspond to regions of high LCF. This is unsurprising since these isolated ridges must have strong longitudinal coupling in order to preserve conduction along the isolated fibre. For those null distributions which are significantly different to the uniform case, we note that the effect sizes are medium or smaller. As expected when comparing the fibre and fibre-less null models, we find statistically significant differences between the LCF distributions, particularly at medium and high risk.

Overall, the effect sizes for the LCF are generally smaller than the equivalent OVF effect sizes. This indicates that OVF is the dominant factor in determining the micro-reentrant substrate in the spatial networks.

Distribution X	Distribution Y	$\lvert\bar{X} - \bar{Y}\rvert$	$\lvert med(X) - med(Y)\rvert$	p-value	Effect size	Descriptor
Fibre Low Risk	Uniform	0.227	0.290	0	2.49	Huge
Fibre Medium Risk	Uniform	0.011	0.039	0.002	0.26	Small / Medium
Fibre High Risk	Uniform	0.006	0.019	0.06	0.15	Small
Null Low Risk	Uniform	0.203	0.272	0	2.20	Huge
Null Medium Risk	Uniform	0.142	0.223	0	2.05	Huge
Null High Risk	Uniform	0.059	0.111	0	0.81	Large
Fibre Low Risk	Null Low Risk	0.020	0.017	0	0.14	Small
Fibre Medium Risk	Null Medium Risk	0.127	0.164	0	1.09	Large / Very Large
Fibre High Risk	Null High Risk	0.064	0.095	0	0.69	Large

Table 6.6: Occupied voxel fraction statistical tests in the healthy sheep atria using a permutation test. Effect size given by nonparametric Cohen's d. Descriptors for nonparametric Cohen's d from [237].

Distribution X	Distribution Y	$\lvert\bar{X} - \bar{Y}\rvert$	$\lvert med(X) - med(Y)\rvert$	p-value	Effect size	Descriptor
Fibre Low Risk	Uniform	0.183	0.217	0	1.41	Very Large
Fibre Medium Risk	Uniform	0.055	0.081	0	0.44	Medium
Fibre High Risk	Uniform	0.043	0.056	0	0.31	Medium
Null Low Risk	Uniform	0.347	0.402	0	2.36	Huge
Null Medium Risk	Uniform	0.101	0.135	0	0.76	Large
Null High Risk	Uniform	0.055	0.094	0	0.53	Medium
Fibre Low Risk	Null Low Risk	0.170	0.192	0	1.71	Very Large / Huge
Fibre Medium Risk	Null Medium Risk	0.043	0.050	0	0.38	Medium
Fibre High Risk	Null High Risk	0.013	0.021	0.008	0.15	Small

Table 6.7: Occupied voxel fraction statistical tests in the HF sheep atria using a permutation test. Effect size given by nonparametric Cohen's d. Descriptors for nonparametric Cohen's d from [237].

Distribution X	Distribution Y	$\lvert\bar{X} - \bar{Y}\rvert$	$\lvert med(X) - med(Y)\rvert$	p-value	Effect size	Descriptor
Fibre Low Risk	Uniform	0.313	0.336	0	1.95	Huge
Fibre Medium Risk	Uniform	0.200	0.238	0	1.29	Very Large
Fibre High Risk	Uniform	0.117	0.141	0	0.77	Large
Null Low Risk	Uniform	0.240	0.282	0	1.59	Very Large / Huge
Null Medium Risk	Uniform	0.216	0.254	0	1.46	Very Large
Null High Risk	Uniform	0.160	0.190	0	1.02	Large / Very Large
Fibre Low Risk	Null Low Risk	0.071	0.060	0	0.79	Large
Fibre Medium Risk	Null Medium Risk	0.013	0.009	0.001	0.07	Small
Fibre High Risk	Null High Risk	0.041	0.047	0	0.36	Medium

Table 6.5: Occupied voxel fraction statistical tests in the human atria using a permutation test. Effect size given by nonparametric Cohen's d. Descriptors for nonparametric Cohen's d from [237].

Distribution X	Distribution Y	$\lvert\bar{X}-\bar{Y}\rvert$	$\lvert med(X)-med(Y)\rvert$	p-value	Effect size	Descriptor
Fibre Low Risk	Uniform	0.036	0.034	0	0.81	Large
Fibre Medium Risk	Uniform	0.034	0.031	0	0.59	Medium / Large
Fibre High Risk	Uniform	0.021	0.022	0	0.48	Medium
Null Low Risk	Uniform	0.017	0.021	0	0.39	Medium
Null Medium Risk	Uniform	0.014	0.018	0	0.33	Medium
Null High Risk	Uniform	0.002	0.007	0.015	0.13	Small
Fibre Low Risk	Null Low Risk	0.053	0.055	0	1.03	Large / Very Large
Fibre Medium Risk	Null Medium Risk	0.048	0.049	0	0.83	Large
Fibre High Risk	Null High Risk	0.023	0.029	0	0.54	Medium

Table 6.8: Longitudinal connection fraction statistical tests in the human atria using a permutation test. Effect size given by nonparametric Cohen's d. Descriptors for nonparametric Cohen's d from [237].

Distribution X	Distribution Y	$\lvert\bar{X}-\bar{Y}\rvert$	$\lvert med(X)-med(Y)\rvert$	p-value	Effect size	Descriptor
Fibre Low Risk	Uniform	0.058	0.066	0	0.63	Medium / Large
Fibre Medium Risk	Uniform	0.035	0.038	0	0.42	Medium
Fibre High Risk	Uniform	0.026	0.030	0	0.35	Medium
Null Low Risk	Uniform	0.003	0.005	0.018	0.04	Very Small
Null Medium Risk	Uniform	0.006	0.001	0	0.01	Very Small
Null High Risk	Uniform	0.005	0.003	0	0.03	Very Small
Fibre Low Risk	Null Low Risk	0.055	0.071	0	0.59	Medium / Large
Fibre Medium Risk	Null Medium Risk	0.028	0.039	0	0.40	Medium
Fibre High Risk	Null High Risk	0.021	0.032	0	0.38	Medium

Table 6.9: Longitudinal connection fraction statistical tests in the healthy sheep atria using a permutation test. Effect size given by nonparametric Cohen's d. Descriptors for nonparametric Cohen's d from [237].

Distribution X	Distribution Y	$\lvert\bar{X}-\bar{Y}\rvert$	$\lvert med(X)-med(Y)\rvert$	p-value	Effect size	Descriptor
Fibre Low Risk	Uniform	0.011	0.008	0	0.10	Small
Fibre Medium Risk	Uniform	0.032	0.036	0	0.42	Medium
Fibre High Risk	Uniform	0.034	0.038	0	0.43	Medium
Null Low Risk	Uniform	0.021	0.025	0	0.38	Medium
Null Medium Risk	Uniform	0.002	0.001	0.085	0.007	Small
Null High Risk	Uniform	0.002	0.001	0.023	0.010	Small
Fibre Low Risk	Null Low Risk	0.010	0.017	0	0.23	Small / Medium
Fibre Medium Risk	Null Medium Risk	0.030	0.036	0	0.39	Medium
Fibre High Risk	Null High Risk	0.031	0.039	0	0.43	Medium

Table 6.10: Longitudinal connection fraction statistical tests in the HF sheep atria using a permutation test. Effect size given by nonparametric Cohen's d. Descriptors for nonparametric Cohen's d from [237].

Bibliography

[1] M. Falkenberg, J.-H. Lee, S.-i. Amano, K.-i. Ogawa, K. Yano, Y. Miyake, T. S. Evans, and K. Christensen. "Identifying time dependence in network growth". In: *Physical Review Research* 2.2 (2020), p. 023352. DOI: 10.1103/PhysRevResearch.2.023352.

[2] M. Falkenberg. "Heterogeneous node copying from hidden network structure". In: *Communications Physics* 4.1 (2021), p. 200. DOI: 10.1038/s42005-021-00694-1.

[3] M. Falkenberg, A. J. Ford, A. C. Li, R. Lawrence, A. Ciacci, N. S. Peters, and K. Christensen. "Unified mechanism of local drivers in a percolation model of atrial fibrillation". In: *Physical Review E* 100.6 (2019), p. 062406. DOI: 10.1103/PhysRevE.100.062406.

[4] M. Falkenberg, D. Hickey, L. Terrill, A. Ciacci, N. S. Peters, and K. Christensen. "Identifying Potential Re-Entrant Circuit Locations From Atrial Fibre Maps". In: *2019 Computing in Cardiology (CinC)*. IEEE. (2019), pp. 1–4. DOI: 10.22489/CinC.2019.102.

[5] M. Falkenberg, J. A. Coleman, S. Dobson, D. J. Hickey, L. Terrill, A. Ciacci, B. Thomas, N. S. Peters, A. Sau, F. S. Ng, J. Zhao, and K. Christensen. "Identifying locations susceptible to micro-anatomical reentry using a spatial network representation of atrial fibre maps". In: *bioRxiv* (2021). DOI: 10.1101/2021.09.13.460069.

[6] M. Falkenberg, W. Cheng, N. S. Peters, and K. Christensen. "Machine learning methods for locating re-entrant drivers from electrograms in a model of atrial fibrillation". In: *Royal Society Open Science* 5.4 (2018), p. 172434. DOI: 10.1098/rsos.172434.

[7] N. R. Franks, A. Worley, M. Falkenberg, A. B. Sendova-Franks, and K. Christensen. "Digging the optimum pit: antlions, spirals and spontaneous stratification". In: *Proceed-*

ings of the Royal Society B 286.1899 (2019), p. 20190365. DOI: 10.1098/rspb.2019.0365.

[8] A. Ciacci, M. Falkenberg, K. A. Manani, T. S. Evans, N. S. Peters, and K. Christensen. "Understanding the transition from paroxysmal to persistent atrial fibrillation". In: *Physical Review Research* 2.2 (2020), p. 023311. DOI: 10.1103/PhysRevResearch.2.023311.

[9] H. Rajpal, M. Sas, C. Lockwood, R. Joakim, N. S. Peters, and M. Falkenberg. "Interpretable XGBoost Based Classification of 12-lead ECGs Applying Information Theory Measures From Neuroscience". In: *2020 Computing in Cardiology*. IEEE. (2020), pp. 1–4. DOI: 10.22489/CinC.2020.185.

[10] V. Vasiliauskaite and F. E. Rosas. "Understanding complexity via network theory: a gentle introduction". In: *Echo* 2 (2021). DOI: 10.47041/CVTD4629.

[11] D. J. Watts and S. H. Strogatz. "Collective dynamics of 'small-world' networks". In: *Nature* 393.6684 (1998), p. 440. DOI: 10.1038/30918.

[12] D. S. Price. "A general theory of bibliometric and other cumulative advantage processes". In: *Journal of the American Society for Information Science* 27 (1976), pp. 292–306. DOI: 10.1002/asi.4630270505.

[13] A.-L. Barabási and R. Albert. "Emergence of Scaling in Random Networks". In: *Science* 286.5439 (1999), pp. 509–512. DOI: 10.1126/science.286.5439.509.

[14] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, and G. Petri. "Networks beyond pairwise interactions: structure and dynamics". In: *Physics Reports* (2020). DOI: 10.1016/j.physrep.2020.05.004.

[15] F. Battiston, E. Amico, A. Barrat, G. Bianconi, G. Ferraz de Arruda, B. Franceschiello, I. Iacopini, S. Kéfi, V. Latora, Y. Moreno, M. M. Murray, T. P. Peixoto, F. Vaccarino, and G. Petri. "The physics of higher-order interactions in complex systems". In: *Nature Physics* 17.10 (2021), pp. 1093–1098. DOI: 10.1038/s41567-021-01371-4.

[16] M. M. Danziger, A. Bashan, Y. Berezin, L. M. Shekhtman, and S. Havlin. "An introduction to interdependent networks". In: *International Conference on Nonlinear Dynamics of Electronic Systems*. Springer. (2014), pp. 189–202. DOI: 10.1007/978-3-319-08672-9_24.

[17] G. Bianconi. *Multilayer networks: structure and function*. Oxford university press, (2018). DOI: 10.1093/oso/9780198753919.001.0001.

[18] T. Carletti, F. Battiston, G. Cencetti, and D. Fanelli. "Random walks on hypergraphs". In: *Physical Review E* 101.2 (2020), p. 022308. DOI: 10.1103/PhysRevE.101.022308.

[19] A. D. Broido and A. Clauset. "Scale-free networks are rare". In: *Nature Communications* 10.1 (2019), p. 1017. DOI: 10.1038/s41467-019-08746-5.

[20] M. E. J. Newman. "Power laws, Pareto distributions and Zipf's law". In: *Contemporary Physics* 46.5 (2005), pp. 323–351. DOI: 10.1080/00107510500052444.

[21] R. Albert, H. Jeong, and A.-L. Barabási. "Diameter of the World-Wide Web". In: *Nature* 401 (1999), p. 130. DOI: 10.1038/43601.

[22] S. Redner. "How popular is your paper? An empirical study of the citation distribution". In: *The European Physical Journal B-Condensed Matter and Complex Systems* 4.2 (1998), pp. 131–134. DOI: 10.1007/s100510050359.

[23] R. F. i Cancho and R. V. Solé. "The small world of human language". In: *Proceedings of the Royal Society of London B: Biological Sciences* 268.1482 (2001), pp. 2261–2265. DOI: 10.1098/rspb.2001.1800.

[24] F. Liljeros, C. R. Edling, L. A. N. Amaral, H. E. Stanley, and Y. Åberg. "The web of human sexual contacts". In: *Nature* 411.6840 (2001), p. 907. DOI: 10.1038/35082140.

[25] A. Java, X. Song, T. Finin, and B. Tseng. "Why we twitter: understanding microblogging usage and communities". In: *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*. ACM. (2007), pp. 56–65. DOI: 10.1145/1348549.1348556.

[26] R. Cohen and S. Havlin. "Scale-Free Networks Are Ultrasmall". In: *Physical Review Letters* 90 (2003), p. 058701. DOI: 10.1103/PhysRevLett.90.058701.

[27] P. Crucitti, V. Latora, M. Marchiori, and A. Rapisarda. "Efficiency of scale-free networks: error and attack tolerance". In: *Physica A: Statistical Mechanics and its Applications* 320 (2003), pp. 622–642. DOI: 10.1016/S0378-4371(02)01545-5.

[28] P. Holme. "Rare and everywhere: Perspectives on scale-free networks". In: *Nature Communications* 10.1 (2019), p. 1016. DOI: 10.1038/s41467-019-09038-8.

[29] A. Clauset, C. R. Shalizi, and M. E. J. Newman. "Power-law distributions in empirical data". In: *SIAM Review* 51.4 (2009), pp. 661–703. DOI: 10.1137/070710111.

[30] F. Radicchi, S. Fortunato, and C. Castellano. "Universality of citation distributions: Toward an objective measure of scientific impact". In: *Proceedings of the National Academy of Sciences* 105 (2008), pp. 17268–17272. DOI: 10.1073/pnas.0806977105.

[31] G. Lima-Mendez and J. van Helden. "The powerful law of the power law and other myths in network biology". In: *Molecular BioSystems* 5.12 (2009), pp. 1482–1493. DOI: 10.1039/B908681A.

[32] W. Willinger, D. Alderson, and J. C. Doyle. "Mathematics and the internet: A source of enormous confusion and great potential". In: *Notices of the American Mathematical Society* 56.5 (2009), pp. 586–599. URL: ams.org/notices/200905/tx090500586p.pdf.

[33] M. Golosovsky. "Power-law citation distributions are not scale-free". In: *Physical Review E* 96.3 (2017), p. 032306. DOI: 10.1103/PhysRevE.95.012324.

[34] M. P. Stumpf and M. A. Porter. "Critical truths about power laws". In: *Science* 335.6069 (2012), pp. 665–666. DOI: 10.1126/science.1216142.

[35] I. Sendiña-Nadal, M. M. Danziger, Z. Wang, S. Havlin, and S. Boccaletti. "Assortativity and leadership emerge from anti-preferential attachment in heterogeneous networks". In: *Scientific Reports* 6.1 (2016), p. 21297. DOI: 10.1038/srep21297.

[36] K.-I. Goh, E. Oh, H. Jeong, B. Kahng, and D. Kim. "Classification of scale-free networks". In: *Proceedings of the National Academy of Sciences* 99.20 (2002), pp. 12583–12588. DOI: 10.1073/pnas.202301299.

[37] T. House, J. M. Read, L. Danon, and M. J. Keeling. "Testing the hypothesis of preferential attachment in social network formation". In: *EPJ Data Science* 4.1 (2015), p. 13. DOI: 10.1140/epjds/s13688-015-0052-2.

[38] A.-L. Barabási. *Network science*. Cambridge university press, (2016). URL: www.networksciencebook.com.

[39] I. Voitalov, P. van der Hoorn, R. van der Hofstad, and D. Krioukov. "Scale-free networks well done". In: *Physical Review Research* 1 (2019), p. 033034. DOI: 10.1103/PhysRevResearch.1.033034.

[40] M. Gerlach and E. G. Altmann. "Testing Statistical Laws in Complex Systems". In: *Physical Review Letters* 122 (2019), p. 168301. DOI: 10.1103/PhysRevLett.122.168301.

[41] I. Z. Kiss, D. M. Green, and R. R. Kao. "The effect of network mixing patterns on epidemic dynamics and the efficacy of disease contact tracing". In: *Journal of The Royal Society Interface* 5.24 (2008), pp. 791–799. DOI: 10.1098/rsif.2007.1272.

[42] M. Piraveenan, M. Prokopenko, and A. Zomaya. "Assortative mixing in directed biological networks". In: *IEEE/ACM Transactions on Computational Biology and Bioinformatics* 9.1 (2012), pp. 66–78. DOI: 10.1109/TCBB.2010.80.

[43] P. L. Krapivsky, S. Redner, and F. Leyvraz. "Connectivity of Growing Random Networks". In: *Physical Review Letters* 85 (2000), pp. 4629–4632. DOI: 10.1103/PhysRevLett.85.4629.

[44] P. Krapivsky and D. Krioukov. "Scale-free networks as preasymptotic regimes of superlinear preferential attachment". In: *Physical Review E* 78.2 (2008), p. 026114. DOI: 10.1103/PhysRevE.78.026114.

[45] M. E. J. Newman. "Clustering and preferential attachment in growing networks". In: *Physical Review E* 64.2 (2001), p. 025102. DOI: 10.1103/PhysRevE.64.025102.

[46] H. Jeong, Z. Néda, and A.-L. Barabási. "Measuring preferential attachment in evolving networks". In: *Europhysics Letters* 61.4 (2003), p. 567. DOI: 10.1209/epl/i2003-00166-9.

[47] C. P. Massen and J. P. K. Doye. "Preferential attachment during the evolution of a potential energy landscape". In: *The Journal of Chemical Physics* 127.11 (2007), p. 114306. DOI: 10.1063/1.2773721.

[48] V. Gómez, H. J. Kappen, and A. Kaltenbrunner. "Modeling the Structure and Evolution of Discussion Cascades". In: *Proceedings of the 22Nd ACM Conference on Hypertext and Hypermedia*. HT '11. Eindhoven, The Netherlands: ACM, (2011), pp. 181–190. ISBN: 978-1-4503-0256-2. DOI: 10.1145/1995966.1995992.

[49] P. Sheridan, Y. Yagahara, and H. Shimodaira. "Measuring preferential attachment in growing networks with missing-timelines using Markov chain Monte Carlo". In: *Physica A: Statistical Mechanics and its Applications* 391.20 (2012), pp. 5031–5040. DOI: 10.1016/j.physa.2012.05.041.

[50] J. Kunegis, M. Blattner, and C. Moser. "Preferential attachment in online networks: Measurement and explanations". In: *Proceedings of the 5th annual ACM web science conference*. ACM. (2013), pp. 205–214. DOI: 10.1145/2464464.2464514.

[51] T. Pham, P. Sheridan, and H. Shimodaira. "PAFit: A statistical method for measuring preferential attachment in temporal complex networks". In: *PloS One* 10.9 (2015), e0137796. DOI: 10.1371/journal.pone.0137796.

[52] A. Herdağdelen, E. Aygün, and H. Bingol. "A formal treatment of generalized preferential attachment and its empirical validation". In: *Europhysics Letters* 78.6 (2007), p. 60007. DOI: 10.1209/0295-5075/78/60007.

[53] J. Leskovec, J. Kleinberg, and C. Faloutsos. "Graphs over time: densification laws, shrinking diameters and possible explanations". In: *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. (2005), pp. 177–187. DOI: 10.1145/1081870.1081893.

[54] G. A. Ronda-Pupo and T. Pham. "The evolutions of the rich get richer and the fit get richer phenomena in scholarly networks: the case of the strategic management journal". In: *Scientometrics* 116.1 (2018), pp. 363–383. DOI: 10.1007/s11192-018-2761-3.

[55] P. Sheridan and T. Onodera. "A Preferential Attachment Paradox: How Preferential Attachment Combines with Growth to Produce Networks with Log-Normal In-Degree Distributions". In: *Scientific Reports* 8.1 (2018), p. 2811. DOI: 10.1038/s41598-018-21133-2.

[56] E. Eisenberg and E. Y. Levanon. "Preferential attachment in the protein network evolution". In: *Physical Review Letters* 91.13 (2003), p. 138701. DOI: 10.1103/PhysRevLett.91.138701.

[57] D. Kondor, M. Pósfai, I. Csabai, and G. Vattay. "Do the rich get richer? An empirical analysis of the Bitcoin transaction network". In: *PloS One* 9.2 (2014), e86197. DOI: 10.1371/journal.pone.0086197.

[58] M. Perc. "Evolution of the most common English words and phrases over the centuries". In: *Journal of The Royal Society Interface* 9.77 (2012), pp. 3323–3328. DOI: 10.1098/rsif.2012.0491.

[59] M. Szell and S. Thurner. "Measuring social dynamics in a massive multiplayer online game". In: *Social Networks* 32.4 (2010), pp. 313–329. DOI: 10.1016/j.socnet.2010.06.001.

[60] W. Li, T. Aste, F. Caccioli, and G. Livan. "Early coauthorship with top scientists predicts success in academic careers". In: *Nature Communications* 10.1 (2019), p. 5170. DOI: 10.1038/s41467-019-13130-4.

[61] A. Vazquez. *Knowing a network by walking on it: emergence of scaling.* (2000). arXiv: cond-mat/0006132 [cond-mat.stat-mech].

[62] R. Lambiotte, P. L. Krapivsky, U. Bhat, and S. Redner. "Structural Transitions in Densifying Networks". In: *Physical Review Letters* 117 (2016), p. 218301. DOI: 10.1103/PhysRevLett.117.218301.

[63] U. Bhat, P. L. Krapivsky, R. Lambiotte, and S. Redner. "Densification and structural transitions in networks that grow by node copying". In: *Physical Review E* 94 (2016), p. 062302. DOI: 10.1103/PhysRevE.94.062302.

[64] P. L. Krapivsky and S. Redner. "Network growth by copying". In: *Physical Review E* 71 (2005), p. 036118. DOI: 10.1103/PhysRevE.71.036118.

[65] C. Dangalchev. "Generation models for scale-free networks". In: *Physica A: Statistical Mechanics and its Applications* 338.3 (2004), pp. 659–671. DOI: 10.1016/j.physa.2004.01.056.

[66] S. Wang, X. Wang, Q. Song, and Y. Zhang. "High-Order Degree and Combined Degree in Complex Networks". In: *Mathematical Problems in Engineering* 2018 (2018). Ed. by E. Zumpano, p. 4925841. DOI: 10.1155/2018/4925841.

[67] A. Topirceanu, M. Udrescu, and R. Marculescu. "Weighted Betweenness Preferential Attachment: A New Mechanism Explaining Social Network Formation and Evolution". In: *Scientific Reports* 8.1 (2018), p. 10871. DOI: 10.1038/s41598-018-29224-w.

[68] S.-i. Amano, K.-i. Ogawa, and Y. Miyake. "Node property of weighted networks considering connectability to nodes within two degrees of separation". In: *Scientific Reports* 8.1 (2018), p. 8464. DOI: 10.1038/s41598-018-26781-y.

[69] A. E. Mislove. "Online social networks: measurement, analysis, and applications to distributed information systems". PhD thesis. Rice University, (2009). DOI: 1911/61861.

[70] M. S. Granovetter. "The strength of weak ties". In: *Social Networks*. Elsevier, (1977), pp. 347–367. DOI: 10.1016/B978-0-12-442450-0.50025-0.

[71] R. S. Burt. *Structural holes: The social structure of competition*. Harvard University Press, (1992). ISBN: 978-0-674-84372-1.

[72] L. Katz. "A New Status Index Derived from Sociometric Index". In: *Psychometrika* 18 (1953), pp. 39–43. DOI: 10.1007\%2FBF02289026.

[73] M. E. J. Newman. *Networks*. 2nd ed. OUP Oxford, (2018). ISBN: 9780192527493. DOI: 10.1093/acprof:oso/9780199206650.001.0001.

[74] S. Brin and L. Page. "The anatomy of a large-scale hypertextual Web search engine". In: *Computer networks and ISDN systems* 30.1-7 (1998), pp. 107–117. DOI: 10.1016/S0169-7552(98)00110-X.

[75] R. Toivonen, L. Kovanen, M. Kivelä, J.-P. Onnela, J. Saramäki, and K. Kaski. "A comparative study of social network models: Network evolution models and nodal attribute models". In: *Social Networks* 31.4 (2009), pp. 240–254. DOI: 10.1016/j.socnet.2009.06.004.

[76] S. R. Goldberg, H. Anthony, and T. S. Evans. "Modelling citation networks". In: *Scientometrics* 105.3 (2015), pp. 1577–1604. DOI: 10.1007/s11192-015-1737-9.

[77] S. Ohno. *Evolution by gene duplication*. Springer Science & Business Media, (2013). DOI: 10.1007/978-3-642-86659-3.

[78] I. Ispolatov, P. L. Krapivsky, and A. Yuryev. "Duplication-divergence model of protein interaction network". In: *Physical Review E* 71.6 (2005), p. 061911. DOI: 10.1103/PhysRevE.71.061911.

[79] J. Kim, P. Krapivsky, B. Kahng, and S. Redner. "Infinite-order percolation and giant fluctuations in a protein interaction network". In: *Physical Review E* 66.5 (2002), p. 055101. DOI: 10.1103/PhysRevE.66.055101.

[80] B. Viswanath, A. Mislove, M. Cha, and K. P. Gummadi. "On the evolution of user interaction in facebook". In: *Proceedings of the 2nd ACM workshop on Online social networks*. (2009), pp. 37–42. DOI: 10.1145/1592665.1592675.

[81] American Physical Society. *APS Data Sets for Research*. journals.aps.org/datasets. (2018).

[82] J. Kunegis. "KONECT: The Koblenz Network Collection". In: *Proceedings of the 22nd International Conference on World Wide Web*. WWW '13 Companion. Rio de Janeiro, Brazil: Association for Computing Machinery, (2013), pp. 1343–1350. ISBN: 9781450320382. DOI: 10.1145/2487788.2488173.

[83] A.-L. Barabási, H. Jeong, Z. Néda, E. Ravasz, A. Schubert, and T. Vicsek. "Evolution of the social network of scientific collaborations". In: *Physica A: Statistical mechanics and its applications* 311.3-4 (2002), pp. 590–614. DOI: 10.1016/S0378-4371(02)00736-7.

[84] C. Knappett, T. Evans, and R. Rivers. "Modelling maritime interaction in the Aegean Bronze Age". In: *Antiquity* 82.318 (2008), pp. 1009–1024. DOI: 10.1017/S0003598X0009774X.

[85] D. S. Bassett and O. Sporns. "Network Neuroscience". In: *Nature Neuroscience* 20.3 (2017), p. 353. DOI: 10.1038/nn.4502.

[86] F. Schweitzer, G. Fagiolo, D. Sornette, F. Vega-Redondo, A. Vespignani, and D. R. White. "Economic networks: The new challenges". In: *Science* 325.5939 (2009), pp. 422–425. DOI: 10.1126/science.1173644.

[87] R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani. "Epidemic processes in complex networks". In: *Reviews of Modern Physics* 87.3 (2015), p. 925. DOI: 10.1103/RevModPhys.87.925.

[88] P. L. Krapivsky and S. Redner. "Organization of growing random networks". In: *Physical Review E* 63 (2001), p. 066123. DOI: 10.1103/PhysRevE.63.066123.

[89] A. Vázquez. "Growing networks with local rules: preferential attachment, clustering hierarchy and degree correlations". In: *Physical Review E* 67 (2003), p. 056104. DOI: 10.1103/PhysRevE.67.056104.

[90] F. Chung, L. Lu, T. G. Dewey, and D. J. Galas. "Duplication models for biological networks". In: *Journal of Computational Biology* 10.5 (2003), pp. 677–687. DOI: 10.1089/106652703322539024.

[91] J. Saramäki and K. Kaski. "Scale-Free Networks Generated by Random Walkers". In: *Physica A: Statistical Mechanics and its Applications* 341 (2004), p. 80. DOI: 10.1016/j.physa.2004.04.110.

[92] T. S. Evans and J. P. Saramäki. "Scale Free Networks from Self-Organisation". In: *Physical Review E* 72 (2005), p. 026138. DOI: 10.1103/PhysRevE.72.026138.

[93] J. Sun, M. Medo, and S. Staab. "Time-invariant degree growth in preferential attachment network models". In: *Physical Review E* 101 (2020), p. 022309. DOI: 10.1103/PhysRevE.101.022309.

[94] V. Amati, A. Lomi, and D. Mascia. "Some days are better than others: Examining time-specific variation in the structuring of interorganizational relations". In: *Social Networks* 57 (2019), pp. 18–33. DOI: 10.1016/j.socnet.2018.10.001.

[95] M. E. J. Newman. "The structure and function of complex networks". In: *SIAM Review* 45.2 (2003), pp. 167–256. DOI: 10.1137/S003614450342480.

[96] U. Bhat, P. Krapivsky, R. Lambiotte, and S. Redner. "Densification and structural transitions in networks that grow by node copying". In: *Physical Review E* 94.6 (2016), p. 062302. DOI: 10.1103/PhysRevE.94.062302.

[97] U. Bhat, P. L. Krapivsky, and S. Redner. "Emergence of clustering in an acquaintance model without homophily". In: *Journal of Statistical Mechanics: Theory and Experiment* 2014.11 (2014), P11035. DOI: 10.1088/1742-5468/2014/11/P11035.

[98] G. Bianconi, R. K. Darst, J. Iacovacci, and S. Fortunato. "Triadic closure as a basic generating mechanism of communities in complex networks". In: *Physical Review E* 90.4 (2014), p. 042806. DOI: 10.1103/PhysRevE.90.042806.

[99] J. Davidsen, H. Ebel, and S. Bornholdt. "Emergence of a small world from local interactions: Modeling acquaintance networks". In: *Physical Review Letters* 88.12 (2002), p. 128701. DOI: 10.1103/PhysRevLett.88.128701.

[100] M. K. Hassan, L. Islam, and S. A. Haque. "Degree distribution, rank-size distribution, and leadership persistence in mediation-driven attachment networks". In: *Physica A: Statistical Mechanics and its Applications* 469 (2017), pp. 23–30. DOI: 10.1016/j.physa.2016.11.001.

[101] R. Toivonen, J.-P. Onnela, J. Saramäki, J. Hyvönen, and K. Kaski. "A model for social networks". In: *Physica A: Statistical Mechanics and its Applications* 371.2 (2006), pp. 851–860. DOI: 10.1016/j.physa.2006.03.050.

[102] A. Asikainen, G. Iñiguez, J. Ureña-Carrión, K. Kaski, and M. Kivelä. "Cumulative effects of triadic closure and homophily in social networks". In: *Science Advances* 6.19 (2020), eaax7310. DOI: 10.1126/sciadv.aax7310.

[103] N. Farid and K. Christensen. "Evolving networks through deletion and duplication". In: *New Journal of Physics* 8.9 (2006), p. 212. DOI: 10.1088/1367-2630/8/9/212.

[104] R. Pastor-Satorras, E. Smith, and R. V. Solé. "Evolving protein interaction networks through gene duplication". In: *Journal of Theoretical Biology* 222.2 (2003), pp. 199–210. DOI: 10.1016/S0022-5193(03)00028-6.

[105] I. Ispolatov, P. L. Krapivsky, and A. Yuryev. "Duplication-divergence model of protein interaction network". In: *Physical Review E* 71.6 (2005), p. 061911. DOI: 10.1103/PhysRevE.71.061911.

[106] A. Bhan, D. J. Galas, and T. G. Dewey. "A duplication growth model of gene expression networks". In: *Bioinformatics* 18.11 (2002), pp. 1486–1493. DOI: 10.1093/bioinformatics/18.11.1486.

[107] P. L. Krapivsky and S. Redner. "Network growth by copying". In: *Physical Review E* 71.3 (2005), p. 036118. DOI: 10.1103/PhysRevE.71.036118.

[108] C. Steinbock, O. Biham, and E. Katzav. "Distribution of shortest path lengths in a class of node duplication network models". In: *Physical Review E* 96.3 (2017), p. 032301. DOI: 10.1103/PhysRevE.96.032301.

[109] A. Vázquez, A. Flammini, A. Maritan, and A. Vespignani. "Modeling of protein interaction networks". In: *Complexus* 1.1 (2003), pp. 38–44. DOI: 10.1159/000067642.

[110] P. Holme and B. J. Kim. "Growing scale-free networks with tunable clustering". In: *Physical Review E* 65.2 (2002), p. 026107. DOI: 10.1103/PhysRevE.65.026107.

[111] A. Vázquez. "Growing network with local rules: Preferential attachment, clustering hierarchy, and degree correlations". In: *Physical Review E* 67.5 (2003), p. 056104. DOI: 10.1103/PhysRevE.67.056104.

[112] T. P. Peixoto. "Disentangling homophily, community structure and triadic closure in networks". In: *arXiv preprint arXiv:2101.02510* (2021). URL: arxiv.org/abs/2101.02510.

[113] F. Battiston, J. Iacovacci, V. Nicosia, G. Bianconi, and V. Latora. "Emergence of multiplex communities in collaboration networks". In: *PloS One* 11.1 (2016), e0147451. DOI: 10.1371/journal.pone.0147451.

[114] S. Li, K. P. Choi, and T. Wu. "Degree distribution of large networks generated by the partial duplication model". In: *Theoretical Computer Science* 476 (2013), pp. 94–108. DOI: 10.1016/j.tcs.2012.12.045.

[115] G. Bebek, P. Berenbrink, C. Cooper, T. Friedetzky, J. Nadeau, and S. C. Sahinalp. "The degree distribution of the generalized duplication model". In: *Theoretical Computer Science* 369.1-3 (2006), pp. 239–249. DOI: 10.1016/j.tcs.2006.08.045.

[116] C. Steinbock, O. Biham, and E. Katzav. "Analytical results for the distribution of shortest path lengths in directed random networks that grow by node duplication". In: *The European Physical Journal B* 92.6 (2019), pp. 1–16. DOI: 10.1140/epjb/e2019-100139-5.

[117] C. Steinbock, O. Biham, and E. Katzav. "Analytical results for the in-degree and out-degree distributions of directed random networks that grow by node duplication". In: *Journal of Statistical Mechanics: Theory and Experiment* 2019.8 (2019), p. 083403. DOI: 10.1088/1742-5468/ab3191.

[118] R. I. M. Dunbar. "The social brain hypothesis". In: *Evolutionary Anthropology: Issues, News, and Reviews* 6.5 (1998), pp. 178–190. DOI: 10.1002/(SICI)1520-6505(1998)6:5<178::AID-EVAN5>3.0.CO;2-8.

[119] P. MacCarron, K. Kaski, and R. Dunbar. "Calling Dunbar's numbers". In: *Social Networks* 47 (2016), pp. 151–155. DOI: 10.1016/j.socnet.2016.06.003.

[120] C. R. McClain. "Practices and promises of Facebook for science outreach: Becoming a "Nerd of Trust"". In: *PLoS Biology* 15.6 (2017), e2002020. DOI: 10.1371/journal.pbio.2002020.

[121] D. Easley and J. Kleinberg. "Networks, crowds, and markets". In: vol. 8. Cambridge university press, (2010). Chap. Strong and Weak Ties. DOI: 10.1017/S0266466609990685.

[122] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási. "Structure and tie strengths in mobile communication networks". In: *Proceedings of the National Academy of Sciences* 104.18 (2007), pp. 7332–7336. DOI: 10.1073/pnas.0610245104.

[123] C. Scholz, M. Atzmueller, M. Kibanov, and G. Stumme. "Predictability of evolving contacts and triadic closure in human face-to-face proximity networks". In: *Social Network Analysis and Mining* 4.1 (2014), p. 217. DOI: 10.1007/s13278-014-0217-1.

[124] J. Kim and J. Diesner. "Over-time measurement of triadic closure in coauthorship networks". In: *Social Network Analysis and Mining* 7.1 (2017), p. 9. DOI: 10.1007/s13278-017-0428-3.

[125] L. Zhou, Y. Yang, X. Ren, F. Wu, and Y. Zhuang. "Dynamic network embedding by modeling triadic closure process". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 32. 1. (2018). URL: yangy.org/works/dynamictriad/dynamic_triad.pdf.

[126] T. Raducha, B. Min, and M. San Miguel. "Coevolving nonlinear voter model with triadic closure". In: *Europhysics Letters* 124.3 (2018), p. 30001. DOI: 10.1209/0295-5075/124/30001.

[127] D. Cellai and G. Bianconi. "Multiplex networks with heterogeneous activities of the nodes". In: *Physical Review E* 93.3 (2016), p. 032302. DOI: 10.1103/PhysRevE.93.032302.

[128] J. Gao, S. V. Buldyrev, H. E. Stanley, and S. Havlin. "Networks formed from interdependent networks". In: *Nature Physics* 8.1 (2012), pp. 40–48. DOI: 10.1038/nphys2180.

[129] P. Erdős and A. Rényi. "On random graphs." In: *Publicationes Mathematicae* 6 (1959), pp. 290–297.

[130] R. A. Rossi and N. K. Ahmed. "The Network Data Repository with Interactive Graph Analytics and Visualization". In: *AAAI*. (2015). URL: networkrepository.com.

[131] M. Bardoscia, P. Barucca, S. Battiston, F. Caccioli, G. Cimini, D. Garlaschelli, F. Saracco, T. Squartini, and G. Caldarelli. "The physics of financial networks". In: *Nature Reviews Physics* (2021), pp. 1–18. DOI: 10.1038/s42254-021-00322-5.

[132] J. Wachs and J. Kertész. "A network approach to cartel detection in public auction markets". In: *Scientific Reports* 9.1 (2019), pp. 1–10. DOI: 10.1038/s41598-019-47198-1.

[133] S. Kéfi, E. L. Berlow, E. A. Wieters, L. N. Joppa, S. A. Wood, U. Brose, and S. A. Navarrete. "Network structure beyond food webs: mapping non-trophic and trophic interactions on Chilean rocky shores". In: *Ecology* 96.1 (2015), pp. 291–303. DOI: 10.1890/13-1424.1.

[134] S. Pilosof, M. A. Porter, M. Pascual, and S. Kéfi. "The multilayer nature of ecological networks". In: *Nature Ecology & Evolution* 1.4 (2017), pp. 1–9. DOI: 10.1038/s41559-017-0101.

[135] I. Iacopini, B. Schäfer, E. Arcaute, C. Beck, and V. Latora. "Multilayer modeling of adoption dynamics in energy demand management". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30.1 (2020), p. 013153. DOI: 10.1063/1.5122313.

[136] S. Majhi, M. Perc, and D. Ghosh. "Chimera states in a multilayer network of coupled and uncoupled neurons". In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 27.7 (2017), p. 073109. DOI: 10.1063/1.4993836.

[137] A. G. Kléber and Y. Rudy. "Basic mechanisms of cardiac impulse propagation and associated arrhythmias". In: *Physiological reviews* 84.2 (2004), pp. 431–488. DOI: 10.1152/physrev.00025.2003.

[138] S. Nattel. "New ideas about atrial fibrillation 50 years on." In: *Nature* 415.6868 (2002), pp. 219–226. DOI: 10.1038/415219a.

[139] S. Nattel and D. Dobrev. "Controversies about atrial fibrillation mechanisms: aiming for order in chaos and whether it matters". In: *Circulation Research* 120.9 (2017), pp. 1396–1398. DOI: 10.1161/CIRCRESAHA.116.310489.

[140] F. S. Ng, B. S. Handa, X. Li, and N. S. Peters. "Toward Mechanism-Directed Electrophenotype Based Treatments for Atrial Fibrillation". In: *Frontiers in Physiology* 11 (2020), p. 987. DOI: 10.3389/fphys.2020.00987.

[141] B. S. Handa, X. Li, N. Baxan, C. H. Roney, A. Shchendrygina, C. A. Mansfield, R. J. Jabbour, D. S. Pitcher, R. A. Chowdhury, N. S. Peters, and F. S. Ng. "Ventricular fibrillation mechanism and global fibrillatory organization are determined by gap junction coupling and fibrosis pattern". In: *Cardiovascular Research* (2020). DOI: 10.1093/cvr/cvaa141.

[142] H. Calkins, G. Hindricks, R. Cappato, Y.-H. Kim, E. B. Saad, L. Aguinaga, J. G. Akar, V. Badhwar, J. Brugada, J. Camm, *et al.* "2017 HRS/EHRA/ECAS/APHRS/ SOLAECE expert consensus statement on catheter and surgical ablation of atrial fibrillation". In: *EP Europace* 20 (2017), e1–160. DOI: 10.1093/europace/eux274.

[143] M. Haïssaguerre, P. Jaïs, D. C. Shah, A. Takahashi, M. Hocini, G. Quiniou, S. Garrigue, A. Le Mouroux, P. Le Métayer, and J. Clémenty. "Spontaneous Initiation of Atrial Fibrillation by Ectopic Beats Originating in the Pulmonary Veins". In: *New England Journal of Medicine* 339.10 (1998), pp. 659–666. DOI: 10.1056/NEJM199809033391003.

[144] B. J. Hansen, J. Zhao, T. A. Csepe, B. T. Moore, N. Li, L. A. Jayne, A. Kalyanasundaram, P. Lim, A. Bratasz, K. A. Powell, *et al.* "Atrial fibrillation driven by micro-anatomic intramural re-entry revealed by simultaneous sub-epicardial and sub-endocardial optical mapping in explanted human hearts". In: *European Heart Journal* 36.35 (2015), pp. 2390–2401. DOI: 10.1093/eurheartj/ehv233.

[145] K. Christensen, K. A. Manani, and N. S. Peters. "Simple Model for Identifying Critical Regions in Atrial Fibrillation". In: *Physical Review Letters* 114 (2015), p. 028104. DOI: 10.1103/PhysRevLett.114.028104.

[146] N. J. Patel, V. Atti, R. D. Mitrani, J. F. Viles-Gonzalez, and J. J. Goldberger. "Global rising trends of atrial fibrillation: a major public health concern". In: *Heart* 104.24 (2018), pp. 1989–1990. DOI: 10.1136/heartjnl-2018-313350.

[147] U. Schotten, D. Dobrev, P. G. Platonov, H. Kottkamp, and G. Hindricks. "Current controversies in determining the main mechanisms of atrial fibrillation". In: *Journal of Internal Medicine* 279.5 (2016), pp. 428–438. DOI: 10.1111/joim.12492.

[148] V. V. Fedorov and B. J. Hansen. "Now You See a Rotor, Now You Don't". In: *JACC: Clinical Electrophysiology* 4.1 (2018), pp. 84–86. DOI: 10.1016/j.jacep.2017.08.014.

[149] I. Mann, B. Sandler, N. Linton, and P. Kanagaratnam. "Drivers of Atrial Fibrillation: Theoretical Considerations and Practical Concerns". In: *Arrhythmia & Electrophysiology Review* 7 (2018), pp. 49–54. DOI: 10.15420/aer.2017.40.3.

[150] A. Verma, C.-y. Jiang, T. R. Betts, J. Chen, I. Deisenhofer, R. Mantovan, L. Macle, C. A. Morillo, W. Haverkamp, R. Weerasooriya, *et al.* "Approaches to Catheter Ablation for Persistent Atrial Fibrillation". In: *New England Journal of Medicine* 372.19 (2015), pp. 1812–1822. DOI: 10.1056/NEJMoa1408288.

[151] L. Macle, P. Khairy, R. Weerasooriya, P. Novak, A. Verma, S. Willems, T. Arentz, I. Deisenhofer, G. Veenhuyzen, C. Scavée, *et al.* "Adenosine-guided pulmonary vein isolation for the treatment of paroxysmal atrial fibrillation: an international, multicentre, randomised superiority trial". In: *Lancet* 386.9994 (2015), pp. 672–679. DOI: 10.1016/S0140-6736(15)60026-5.

[152] S. M. Narayan, M. N. Vishwanathan, C. A. B. Kowalewski, T. Baykaner, M. Rodrigo, J. A. B. Zaman, and P. J. Wang. "The continuous challenge of AF ablation: From foci to rotational activity". In: *Revista Portuguesa de Cardiologia* 36 (2017), pp. 9–17. DOI: 10.1016/j.repc.2017.09.007.

[153] K. Shivkumar, K. A. Ellenbogen, J. D. Hummel, J. M. Miller, and J. S. Steinberg. "Acute termination of human atrial fibrillation by identification and catheter ablation of localized rotors and sources: first multicenter experience of focal impulse and rotor modulation (FIRM) ablation". In: *Journal of Cardiovascular Electrophysiology* 23.12 (2012), pp. 1277–1285. DOI: 10.1111/jce.12000.

[154] S. Lee, J. Sahadevan, C. M. Khrestian, I. Cakulev, A. Markowitz, and A. L. Waldo. "Simultaneous Bi-Atrial High Density (510 - 512 electrodes) Epicardial Mapping of

Persistent and Long-Standing Persistent Atrial Fibrillation in Patients: New Insights into the Mechanism of Its Maintenance". In: *Circulation* 132 (2015), pp. 2108–2117. DOI: 10.1161/CIRCULATIONAHA.115.017007.

[155] S. M. Narayan, D. E. Krummen, P. Clopton, K. Shivkumar, and J. M. Miller. "Direct or coincidental elimination of stable rotors or focal sources may explain successful atrial fibrillation ablation". In: *Journal of the American College of Cardiology* 62.2 (2013), pp. 138–147. DOI: 10.1016/j.jacc.2013.03.021.

[156] M. Haïssaguerre, M. Hocini, A. Denis, A. J. Shah, Y. Komatsu, S. Yamashita, M. Daly, S. Amraoui, S. Zellerhoff, M.-Q. Picat, *et al.* "Driver domains in persistent atrial fibrillation". In: *Circulation* 130.7 (2014), pp. 530–538. DOI: 10.1161/CIRCULATIONAHA.113.005421.

[157] B. J. Hansen, T. A. Csepe, J. Zhao, A. J. Ignozzi, J. D. Hummel, and V. V. Fedorov. "Maintenance of Atrial Fibrillation: Are Reentrant Drivers with Spatial Stability the Key?" In: *Circulation: Arrhythmia and Electrophysiology* 9.10 (2016), e004398. DOI: 10.1161/CIRCEP.116.004398.

[158] T. A. Csepe, B. J. Hansen, and V. V. Fedorov. "Atrial fibrillation driver mechanisms: Insight from the isolated human heart". In: *Trends in Cardiovascular Medicine* 27.1 (2017), pp. 1–11. DOI: 10.1016/j.tcm.2016.05.008.

[159] T. Baykaner, A. J. Rogers, G. L. Meckler, J. A. B. Zaman, R. Navara, M. Rodrigo, M. Alhusseini, C. A. B. Kowalewski, M. N. Viswanathan, S. M. Narayan, *et al.* "Clinical Implications of Ablation of Drivers for Atrial Fibrillation: A Systematic Review and Meta-Analysis". In: *Circulation: Arrhythmia and Electrophysiology* 11.5 (2018), e006119. DOI: 10.1161/CIRCEP.117.006119.

[160] J. A. B. Zaman, W. H. Sauer, M. I. Alhusseini, T. Baykaner, R. T. Borne, C. A. B. Kowalewski, S. Busch, P. C. Zei, S. Park, M. N. Viswanathan, *et al.* "Identification and Characterization of Sites Where Persistent Atrial Fibrillation Is Terminated by Localized Ablation." In: *Circulation: Arrhythmia and Electrophysiology* 11.1 (2018), e005258. DOI: 10.1161/CIRCEP.117.005258.

[161] S. Nattel, F. Xiong, and M. Aguilar. "Demystifying rotors and their place in clinical translation of atrial fibrillation mechanisms". In: *Nature Reviews Cardiology* 14.9 (2017), pp. 509–520. DOI: 10.1038/nrcardio.2017.37.

[162] C. X. Wong, A. N. Ganesan, and J. B. Selvanayagam. "Epicardial fat and atrial fibrillation: current evidence, potential mechanisms, clinical implications, and future directions". In: *European Heart Journal* 38.17 (2016), pp. 1294–1302. DOI: 10.1093/eurheartj/ehw045.

[163] C. H. Roney, C. D. Cantwell, J. D. Bayer, N. A. Qureshi, P. B. Lim, J. H. Tweedy, P. Kanagaratnam, N. S. Peters, E. J. Vigmond, and F. S. Ng. "Spatial resolution requirements for accurate identification of drivers of atrial fibrillation". In: *Circulation: Arrhythmia and Electrophysiology* 10.5 (2017), e004899. DOI: 10.1161/CIRCEP.116.004899.

[164] R. H. Clayton, O. Bernus, E. M. Cherry, H. Dierckx, F. H. Fenton, L. Mirabella, A. V. Panfilov, F. B. Sachse, G. Seemann, and H. Zhang. "Models of cardiac tissue electrophysiology: progress, challenges and open questions". In: *Progress in Biophysics and Molecular Biology* 104.1 (2011), pp. 22–48. DOI: 10.1016/j.pbiomolbio.2010.05.008.

[165] N. Wiener and A. Rosenblueth. "The mathematical formulation of the problem of conduction of impulses in a network of connected excitable elements, specifically in cardiac muscle". In: *Archivos del Instituto de Cardiologia de Mexico* 16.3–4 (1946). URL: itlab.us/pubs/wiener_rosenblueth.pdf.

[166] G. K. Moe and J. A. Abildskov. "Atrial fibrillation as a self sustaining arrhythmia independent of focal discharge". In: *American Heart Journal* 58 (1959), pp. 59–70. DOI: 10.1016/0002-8703(59)90274-1.

[167] S. Alonso and M. Bär. "Reentry near the percolation threshold in a heterogeneous discrete model for cardiac tissue". In: *Physical Review Letters* 110.15 (2013), p. 158101. DOI: 10.1103/PhysRevLett.110.158101.

[168] E. Vigmond, A. Pashaei, S. Amraoui, H. Cochet, and M. Haïssaguerre. "Percolation as a mechanism to explain atrial fractionated electrograms and reentry in a fibrosis

model based on imaging data". In: *Heart Rhythm* 13.7 (2016), pp. 1536–1543. DOI: 10.1016/j.hrthm.2016.03.019.

[169] S. Alonso, R. W. dos Santos, and M. Bär. "Reentry and ectopic pacemakers emerge in a three-dimensional model for a slab of cardiac tissue with diffuse microfibrosis near the percolation threshold". In: *PloS one* 11.11 (2016), e0166972. DOI: 10.1371/journal.pone.0166972.

[170] Y. T. Lin, E. T. Y. Chang, J. Eatock, T. Galla, and R. H. Clayton. "Mechanisms of stochastic onset and termination of atrial fibrillation studied with a cellular automaton model". In: *Journal of the Royal Society Interface* 14.128 (2017), p. 20160968. DOI: 10.1098/rsif.2016.0968.

[171] D. Makowiec, J. Wdowczyk, and Z. R. Struzik. "Heart Rhythm Insights Into Structural Remodeling in Atrial Tissue: Timed Automata Approach". In: *Frontiers in Physiology* 9 (2019), p. 1859. DOI: 10.3389/fphys.2018.01859.

[172] T. A. Gokhale, E. Medvescek, and C. S. Henriquez. "Modeling dynamics in diseased cardiac tissue: Impact of model choice". In: *Chaos* 27.9 (2017), p. 093909. DOI: 10.1063/1.4999605.

[173] S. A. Niederer, J. Lumens, and N. A. Trayanova. "Computational models in cardiology". In: *Nature Reviews Cardiology* 16.2 (2019), pp. 100–111. DOI: 10.1038/s41569-018-0104-y.

[174] M. S. Spach and J. P. Boineau. "Microfibrosis Produces Electrical Load Variations Due to Loss of Side-to-Side Cell Connections; A Major Mechanism of Structural Heart Disease Arrhythmias". In: *Pacing and Clinical Electrophysiology* 20.2 (2006), pp. 397–413. DOI: 10.1111/j.1540-8159.1997.tb06199.x.

[175] H. Cochet, R. Dubois, S. Yamashita, N. Al Jefairi, B. Berte, J.-M. Sellal, D. Hooks, A. Frontera, S. Amraoui, A. Zemoura, *et al.* "Relationship Between Fibrosis Detected on Late Gadolinium-Enhanced Cardiac Magnetic Resonance and Re-Entrant Activity Assessed With Electrocardiographic Imaging in Human Persistent Atrial Fibrillation".

In: *JACC: Clinical Electrophysiology* 4 (2018), pp. 17–29. DOI: 10.1016/j.jacep.2017.07.019.

[176] V. Swarup, T. Baykaner, A. Rostamian, J. P. Daubert, J. Hummel, D. E. Krummen, R. Trikha, J. M. Miller, G. F. Tomassoni, and S. M. Narayan. "Stability of rotors and focal sources for human atrial fibrillation: focal impulse and rotor mapping (FIRM) of AF sources and fibrillatory conduction". In: *Journal of Cardiovascular Electrophysiology* 25.12 (2014), pp. 1284–1292. DOI: 10.1111/jce.12559.

[177] R. O. Y. Beinart, S. Abbara, A. Blum, M. Ferencik, K. Heist, J. Ruskin, and M. Mansour. "Left atrial wall thickness variability measured by CT scans in patients undergoing pulmonary vein isolation". In: *Journal of Cardiovascular Electrophysiology* 22.11 (2011), pp. 1232–1236. DOI: 10.1111/j.1540-8167.2011.02100.x.

[178] S. Kumar, C. R. Barbhaiya, S. Balindger, R. M. John, L. M. Epstein, B. A. Koplan, U. B. Tedrow, W. G. Stevenson, and G. F. Michaud. "Better Lesion Creation And Assessment During Catheter Ablation". In: *Journal of Atrial Fibrillation* 8.3 (2015), p. 1189. DOI: 10.4022/jafib.1189.

[179] K. Christensen and N. R. Moloney. *Complexity and Criticality*. Imperial College Press, (2005). DOI: 10.1142/p365.

[180] A. A. Saberi. "Recent advances in percolation theory and its applications". In: *Physics Reports* 578 (2015), pp. 1–32. DOI: 10.1016/j.physrep.2015.03.003.

[181] S. Yen Ho, D. Sanchez-Quintana, J. A. Cabrera, and R. H. Anderson. "Anatomy of the Left Atrium". In: *Journal of Cardiovascular Electrophysiology* 10.11 (1999), pp. 1525–1533. DOI: 10.1111/j.1540-8167.1999.tb00211.x.

[182] C. Weiss, A. Gocht, S. Willems, M. Hoffmann, T. Risius, and T. Meinertz. "Impact of the Distribution and Structure of Myocardium in the Pulmonary Veins for Radiofrequency Ablation of Atrial Fibrillation". In: *Pacing and Clinical Electrophysiology* 25.9 (2002), pp. 1352–1356. DOI: 10.1046/j.1460-9592.2002.01352.x.

[183] R. Weerasooriya, P. Jaïs, C. Scavée, L. Macle, D. C. Shah, T. Arentz, J. A. Salerno, F. Raybaud, K.-J. Choi, M. Hocini, J. Clémenty, and M. Haïssaguerre. "Dissociated Pulmonary Vein Arrhythmia: Incidence and characteristics". In: *Journal of Cardiovascular Electrophysiology* 14.11 (2003), pp. 1173–1179. DOI: 10.1046/j.1540-8167.2003.02583.x.

[184] R. M. Lang, M. Bierig, R. B. Devereux, F. A. Flachskampf, E. Foster, P. A. Pellikka, M. H. Picard, M. J. Roman, J. Seward, J. S. Shanewise, S. D. Solomon, *et al.* "Recommendations for chamber quantification: a report from the American Society of Echocardiography's Guidelines and Standards Committee and the Chamber Quantification Writing Group, developed in conjunction with the European Association of Echocardiograph". In: *Journal of the American Society of Echocardiography* 18.12 (2005), pp. 1440–63. DOI: 10.1016/j.echo.2005.10.005.

[185] A. Maceira, J. Cosín-Sales, M. Roughton, S. Prasad, and D. Pennell. "Reference left atrial dimensions and volumes by steady state free precession cardiovascular magnetic resonance". In: *Journal of Cardiovascular Magnetic Resonance* 12.1 (2010), p. 65. DOI: 10.1186/1532-429X-12-65.

[186] K. Nakamura, N. Funabashi, M. Uehara, M. Ueda, T. Murayama, H. Takaoka, and I. Komuro. "Left atrial wall thickness in paroxysmal atrial fibrillation by multislice-CT is initial marker of structural remodeling and predictor of transition from paroxysmal to chronic form". In: *International Journal of Cardiology* 148.2 (2011), pp. 139–147. DOI: 10.1016/j.ijcard.2009.10.032.

[187] R. A. Luke and J. E. Saffitz. "Remodeling of ventricular conduction pathways in healed canine infarct border zones." In: *Journal of Clinical Investigation* 87 (1991), pp. 1594–602. DOI: 10.1172/JCI115173.

[188] S. Verheule, E. Wilson, T. Everett, S. Shanbhag, C. Golden, and J. Olgin. "Alterations in atrial electrophysiology and tissue structure in a canine model of chronic atrial dilatation due to mitral regurgitation." In: *Circulation* 107 (2003), pp. 2615–22. DOI: 10.1161/01.CIR.0000066915.15187.51.

[189] J. Zhao, U. Schotten, B. Smaill, and S. Verheule. "Loss of Side-to-Side Connections Affects the Relative Contributions of the Sodium and Calcium Current to Transverse Propagation between Strands of Atrial Myocytes". In: *Frontiers in Physiology* 9 (2018), p. 1212. DOI: 10.3389/fphys.2018.01212.

[190] K. S. McDowell, F. Vadakkumpadan, R. Blake, J. Blauer, G. Plank, R. S. MacLeod, and N. A. Trayanova. "Methodology for patient-specific modeling of atrial fibrosis as a substrate for atrial fibrillation". In: *Journal of Electrocardiology* 45.6 (2012), pp. 640–645. DOI: 10.1016/j.jelectrocard.2012.08.005.

[191] P. M. Boyle, T. Zghaib, S. Zahid, R. L. Ali, D. Deng, W. H. Franceschi, J. B. Hakim, M. J. Murphy, A. Prakosa, S. L. Zimmerman, *et al.* "Computationally guided personalized targeted ablation of persistent atrial fibrillation". In: *Nature Biomedical Engineering* 3.11 (2019), pp. 870–879. DOI: 10.1038/s41551-019-0437-9.

[192] J. Dall and M. Christensen. "Random geometric graphs". In: *Physical Review E* 66.1 (2002), p. 016121. DOI: 10.1103/PhysRevE.66.016121.

[193] J. Zhao, B. J. Hansen, Y. Wang, T. A. Csepe, L. V. Sul, A. Tang, Y. Yuan, N. Li, A. Bratasz, K. A. Powell, *et al.* "Three-dimensional integrated functional, structural, and computational mapping to define the structural "fingerprints" of heart-specific atrial fibrillation drivers in human heart ex vivo". In: *Journal of the American Heart Association* 6.8 (2017), e005922. DOI: 10.1161/JAHA.117.005922.

[194] M. Barthélemy. "Spatial networks". In: *Physics Reports* 499.1-3 (2011), pp. 1–101. DOI: 10.1016/j.physrep.2010.11.002.

[195] M. W. Krueger, V. Schmidt, C. Tobón, F. M. Weber, C. Lorenz, D. U. J. Keller, H. Barschdorf, M. Burdumy, P. Neher, G. Plank, *et al.* "Modeling Atrial Fiber Orientation in Patient-Specific Geometries: A Semi-automatic Rule-Based Approach". In: *Functional Imaging and Modeling of the Heart*. Ed. by D. N. Metaxas and L. Axel. Berlin, Heidelberg: Springer Berlin Heidelberg, (2011), pp. 223–232. ISBN: 978-3-642-21028-0. DOI: 10.1007/978-3-642-21028-0_28.

[196] J. Zhao, M. W. Krueger, G. Seemann, S. Meng, H. Zhang, O. Dössel, I. J. LeGrice, and B. H. Smaill. "Myofiber orientation and electrical activation in human and sheep atrial models". In: *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. (2012), pp. 6365–6368. DOI: 10.1109/EMBC.2012.6347450.

[197] J. Zhao, T. D. Butters, H. Zhang, A. J. Pullan, I. J. LeGrice, G. B. Sands, and B. H. Smaill. "An Image-Based Model of Atrial Muscular Architecture: Effects of Structural Anisotropy on Electrical Activation". In: *Circulation: Arrhythmia and Electrophysiology* 5.2 (2012), pp. 361–370. DOI: 10.1161/CIRCEP.111.967950.

[198] B. Thomas. "An imaging pipeline for extracting atrial tissue architecture". PhD thesis. Auckland, New Zealand: University of Auckland, (2020). URL: hdl.handle.net/2292/54052.

[199] J. Bærentzen. *On the implementation of fast marching methods for 3D lattices*. Technical University of Denmark, (2001). URL: www.imm.dtu.dk/pubdb/p.php?841.

[200] S. Y. Ho and D. Sánchez-Quintana. "The importance of atrial structure and fibers". In: *Clinical Anatomy: The Official Journal of the American Association of Clinical Anatomists and the British Association of Clinical Anatomists* 22.1 (2009), pp. 52–63. DOI: 10.1002/ca.20634.

[201] D. Merhof, M. Sonntag, F. Enders, P. Hastreiter, R. Fahlbusch, C. Nimsky, and G. Greiner. "Visualization of diffusion tensor data using evenly spaced streamlines". In: *Vision, Modeling and Visualization* (2005), pp. 257–264. URL: nbn-resolving.de/urn:nbn:de:bsz:352-opus-92130.

[202] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. "Introduction to algorithms". In: MIT press, (2009). DOI: 10.5555/1614191.

[203] J. A. Sethian and A. Vladimirsky. "Fast methods for the Eikonal and related Hamilton–Jacobi equations on unstructured meshes". In: *Proceedings of the National Academy of Sciences* 97.11 (2000), pp. 5699–5703. DOI: 10.1073/pnas.090060097.

[204] A. Loewe, E. Poremba, T. Oesterlein, A. Luik, C. Schmitt, G. Seemann, and O. Dössel. "Patient-specific identification of atrial flutter vulnerability–a computational approach to reveal latent reentry pathways". In: *Frontiers in Physiology* 9 (2019), p. 1910. DOI: 10.3389/fphys.2018.01910.

[205] G. Bub, A. Shrier, and L. Glass. "Spiral wave generation in heterogeneous excitable media". In: *Physical Review Letters* 88.5 (2002), p. 058101. DOI: 10.1103/PhysRevLett.88.058101.

[206] E. J. Ciaccio, N. S. Peters, and H. Garan. "Use of an automaton model to suggest methods for cessation of intractable fibrillatory activity". In: *Computers in Biology and Medicine* 102 (2018), pp. 357–368. DOI: 10.1016/j.compbiomed.2018.07.020.

[207] D. Dharmaprani, M. Schopp, P. Kuklik, D. Chapman, A. Lahiri, L. Dykes, F. Xiong, M. Aguilar, B. Strauss, L. Mitchell, *et al.* "Renewal Theory as a Universal Quantitative Framework to Characterize Phase Singularity Regeneration in Mammalian Cardiac Fibrillation". In: *Circulation: Arrhythmia and Electrophysiology* 12.12 (2019), e007569. DOI: 10.1161/CIRCEP.119.007569.

[208] D. Dharmaprani, E. Jenkins, M. Aguilar, J. X. Quah, A. Lahiri, K. Tiver, L. Mitchell, P. Kuklik, C. Meyer, S. Willems, R. Clayton, M. Nash, S. Nattel, A. D. McGavigan, and A. N. Ganesan. "M/M/Infinity Birth-Death Processes – A Quantitative Representational Framework to Summarize and Explain Phase Singularity and Wavelet Dynamics in Atrial Fibrillation". In: *Frontiers in Physiology* 11 (2021), p. 1786. DOI: 10.3389/fphys.2020.616866.

[209] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. "A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise". In: *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*. KDD'96. Portland, Oregon: AAAI Press, (1996), pp. 226–231. DOI: 10.5555/3001460.3001507.

[210] M. Haïssaguerre, A. J. Shah, H. Cochet, M. Hocini, R. Dubois, I. Efimov, E. Vigmond, O. Bernus, and N. Trayanova. "Intermittent drivers anchoring to structural heterogeneities

as a major pathophysiological mechanism of human persistent atrial fibrillation". In: *The Journal of physiology* 594.9 (2016), pp. 2387–2398. DOI: 10.1113/JP270617.

[211] B. Efron and R. J. Tibshirani. *An introduction to the bootstrap*. CRC press, (1994). DOI: 10.1201/9780429246593.

[212] A. V. Mikhailov, A. Kalyanasundaram, N. Li, S. S. Scott, E. J. Artiga, M. M. Subr, J. Zhao, B. J. Hansen, J. D. Hummel, and V. V. Fedorov. "Comprehensive evaluation of electrophysiological and 3D structural features of human atrial myocardium with insights on atrial fibrillation maintenance mechanisms". In: *Journal of Molecular and Cellular Cardiology* 151 (2021), pp. 56–71. DOI: 10.1016/j.yjmcc.2020.10.012.

[213] B. J. Hansen, J. Zhao, K. M. Helfrich, N. Li, A. Iancau, A. M. Zolotarev, S. O. Zakharkin, A. Kalyanasundaram, M. Subr, N. Dastagir, *et al.* "Unmasking Arrhythmogenic Hubs of Reentry Driving Persistent Atrial Fibrillation for Patient-Specific Treatment". In: *Journal of the American Heart Association* 9 (2020), e017789. DOI: 10.1161/JAHA.120.017789.

[214] H. J. Arevalo, F. Vadakkumpadan, E. Guallar, A. Jebb, P. Malamas, K. C. Wu, and N. A. Trayanova. "Arrhythmia risk stratification of patients after myocardial infarction using personalized heart models". In: *Nature communications* 7.1 (2016), pp. 1–8. DOI: 10.1038/ncomms11437.

[215] C. Sohns and N. F. Marrouche. "Atrial fibrillation and cardiac fibrosis". In: *European Heart Journal* 41.10 (2020), pp. 1123–1131. DOI: 10.1093/eurheartj/ehz786.

[216] R. Sachetto, S. Alonso, and R. W. dos Santos. "Killing Many Birds With Two Stones: Hypoxia and Fibrosis Can Generate Ectopic Beats in a Human Ventricular Model". In: *Frontiers in Physiology* 9 (2018), p. 764. DOI: 10.3389/fphys.2018.00764.

[217] C. H. Roney, J. D. Bayer, S. Zahid, M. Meo, P. M. Boyle, N. A. Trayanova, M. Haïssaguerre, R. Dubois, H. Cochet, and E. J. Vigmond. "Modelling methodology of atrial fibrosis affects rotor dynamics and electrograms". In: *EP Europace* 18 (2016), pp. iv146–iv155. DOI: 10.1093/europace/euw365.

[218] A. Roy, M. Varela, and O. Aslanidi. "Image-based computational evaluation of the effects of atrial wall thickness and fibrosis on re-entrant drivers for atrial fibrillation". In: *Frontiers in Physiology* 9 (2018), p. 1352. DOI: 10.3389/fphys.2018.01352.

[219] S. Kharche, S. Castro, B. Thomas, M. Colman, J. Jarvis, B. Smaill, H. Zhang, R. Stephenson, and J. Zhao. "Role of fiber orientation in atrial arrythmogenesis". In: *Computing in Cardiology 2014*. IEEE. (2014), pp. 1041–1044. URL: cinc.org/archives/2014/pdf/1041.pdf.

[220] C. H. Roney, J. D. Bayer, H. Cochet, M. Meo, R. Dubois, P. Jaïs, and E. J. Vigmond. "Variability in pulmonary vein electrophysiology and fibrosis determines arrhythmia susceptibility and dynamics". In: *PLoS Computational Biology* 14.5 (2018), e1006166. DOI: 10.1371/journal.pcbi.1006166.

[221] C. Roney, R. Bendikas, F. Pashakhanloo, C. Corrado, E. Vigmond, E. McVeigh, N. Trayanova, and S. Niederer. "Constructing a Human Atrial Fibre Atlas." In: *Annals of Biomedical Engineering* (2020). DOI: 10.1007/s10439-020-02525-w.

[222] M. Yamazaki, S. Mironov, C. Taravant, J. Brec, L. M. Vaquero, K. Bandaru, U. M. R. Avula, H. Honjo, I. Kodama, O. Berenfeld, *et al.* "Heterogeneous atrial wall thickness and stretch promote scroll waves anchoring during atrial fibrillation". In: *Cardiovascular Research* 94.1 (2012), pp. 48–57. DOI: 10.1093/cvr/cvs093.

[223] O. V. Aslanidi, M. A. Colman, J. Stott, H. Dobrzynski, M. R. Boyett, A. V. Holden, and H. Zhang. "3D virtual human atria: A computational platform for studying clinical atrial fibrillation". In: *Progress in biophysics and molecular biology* 107.1 (2011), pp. 156–168. DOI: 10.1016/j.pbiomolbio.2011.06.011.

[224] A. Roy, M. Varela, H. Chubb, R. MacLeod, J. C. Hancox, T. Schaeffter, and O. Aslanidi. "Identifying locations of re-entrant drivers from patient-specific distribution of fibrosis in the left atrium". In: *PLoS Computational Biology* 16.9 (2020), e1008086. DOI: 10.1371/journal.pcbi.1008086.

[225] K. N. Aronis, R. Ali, and N. A. Trayanova. "The role of personalized atrial modeling in understanding atrial fibrillation mechanisms and improving treatment". In: *International Journal of Cardiology* 287 (2019), pp. 139–147. DOI: 10.1016/j.ijcard.2019.01.096.

[226] F. Pashakhanloo, D. A. Herzka, H. Ashikaga, S. Mori, N. Gai, D. A. Bluemke, N. A. Trayanova, and E. R. McVeigh. "Myofiber Architecture of the Human Atria as Revealed by Submillimeter Diffusion Tensor Imaging". In: *Circulation: Arrhythmia and Electrophysiology* 9.4 (2016), e004133. DOI: 10.1161/CIRCEP.116.004133.

[227] C. H. Roney, J. Whitaker, I. Sim, L. O'Neill, R. K. Mukherjee, O. Razeghi, E. J. Vigmond, M. Wright, M. D. O'Neill, S. E. Williams, *et al.* "A technique for measuring anisotropy in atrial conduction to estimate conduction velocity and atrial fibre direction". In: *Computers in Biology and Medicine* 104 (2019), pp. 278–290. DOI: 10.1016/j.compbiomed.2018.10.019.

[228] M. W. Krueger, K. S. Rhode, M. D. O'Neill, C. A. Rinaldi, J. Gill, R. Razavi, G. Seemann, and O. Doessel. "Patient-specific modeling of atrial fibrosis increases the accuracy of sinus rhythm simulations and may explain maintenance of atrial fibrillation". In: *Journal of Electrocardiology* 47.3 (2014), pp. 324–328. DOI: 10.1016/j.jelectrocard.2013.11.003.

[229] M. S. Spach and P. C. Dolber. "Relating extracellular potentials and their derivatives to anisotropic propagation at a microscopic level in human cardiac muscle. Evidence for electrical uncoupling of side-to-side fiber connections with increasing age." In: *Circulation Research* 58.3 (1986), pp. 356–371. DOI: 10.1161/01.res.58.3.356.

[230] M. S. Spach, P. Dolber, and J. F. Heidlage. "Influence of the passive anisotropic properties on directional differences in propagation following modification of the sodium conductance in human atrial muscle. A model of reentry based on anisotropic discontinuous propagation." In: *Circulation Research* 62.4 (1988), pp. 811–832. DOI: 10.1161/01.res.62.4.811.

[231] M. W. A. Caan. "DTI Analysis Methods: Fibre Tracking and Connectivity". In: *Diffusion Tensor Imaging: A Practical Handbook*. Ed. by W. Van Hecke, L. Emsell, and

S. Sunaert. New York, NY: Springer New York, (2016), pp. 205–228. ISBN: 978-1-4939-3118-7. DOI: 10.1007/978-1-4939-3118-7_11.

[232] S. Mori, B. J. Crain, V. P. Chacko, and P. C. M. Van Zijl. "Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging". In: *Annals of Neurology* 45.2 (1999), pp. 265–269. DOI: 10.1002/1531-8249(199902)45:2<265::AID-ANA21>3.0.CO;2-3.

[233] Y. Sun, Q. Bi, X. Wang, X. Hu, H. Li, X. Li, T. Ma, J. Lu, P. Chan, N. Shu, and Y. Han. "Prediction of Conversion From Amnestic Mild Cognitive Impairment to Alzheimer's Disease Based on the Brain Structural Connectome". In: *Frontiers in Neurology* 9 (2019), p. 1178. DOI: 10.3389/fneur.2018.01178.

[234] S. Y. Ho, R. H. Anderson, and D. Sánchez-Quintana. "Atrial structure and fibres: morphologic bases of atrial conduction". In: *Cardiovascular Research* 54.2 (2002), pp. 325–336. DOI: 10.1016/S0008-6363(02)00226-2.

[235] Y. Wang, Z. Xiong, A. Nalar, B. J. Hansen, S. Kharche, G. Seemann, A. Loewe, V. V. Fedorov, and J. Zhao. "A robust computational framework for estimating 3D Bi-Atrial chamber wall thickness". In: *Computers in Biology and Medicine* 114 (2019), p. 103444. DOI: 10.1016/j.compbiomed.2019.103444.

[236] R. Wilcox. "A robust nonparametric measure of effect size based on an analog of Cohen's d, plus inferences about the median of the typical difference". In: *Journal of Modern Applied Statistical Methods* 17.2 (2019), p. 1. DOI: 10.22237/jmasm/1551905677.

[237] S. S. Sawilowsky. "New effect size rules of thumb". In: *Journal of Modern Applied Statistical Methods* 8.2 (2009), p. 26. DOI: 10.22237/jmasm/1257035100.

Right to reuse figures

OXFORD UNIVERSITY PRESS LICENSE
TERMS AND CONDITIONS

Sep 20, 2021

This Agreement between Max Falkenberg ("You") and Oxford University Press ("Oxford University Press") consists of your license details and the terms and conditions provided by Oxford University Press and Copyright Clearance Center.

License Number	5141361120934
License date	Sep 03, 2021
Licensed content publisher	Oxford University Press
Licensed content publication	European Heart Journal
Licensed content title	Atrial fibrillation driven by micro-anatomic intramural re-entry revealed by simultaneous sub-epicardial and sub-endocardial optical mapping in explanted human hearts
Licensed content author	Hansen, Brian J.; Zhao, Jichao
Licensed content date	Jun 9, 2015
Type of Use	Thesis/Dissertation
Institution name	
Title of your work	Studies in Complex Systems: Complex Networks With Hidden Layers & Modelling Atrial Fibrillation
Publisher of your work	Imperial College London
Expected publication date	Jan 2022
Permissions cost	0.00 GBP
Value added tax	0.00 GBP
Total	0.00 GBP
Title	Studies in Complex Systems: Complex Networks With Hidden Layers & Modelling Atrial Fibrillation
Institution name	Imperial College London
Expected presentation date	Jan 2022
Portions	Part of Figure 7
Requestor Location	Max Falkenberg Exhibition Road South Kensington London, --- Select One --- SW7 2BX United Kingdom Attn: Max Falkenberg
Publisher Tax ID	GB125506730
Total	0.00 GBP

Terms and Conditions

STANDARD TERMS AND CONDITIONS FOR REPRODUCTION OF MATERIAL FROM AN OXFORD UNIVERSITY PRESS JOURNAL

1. Use of the material is restricted to the type of use specified in your order details.

2. This permission covers the use of the material in the English language in the following territory: world. If you have requested additional permission to translate this material, the terms and conditions of this reuse will be set out in clause 12.

3. This permission is limited to the particular use authorized in (1) above and does not allow you to sanction its use elsewhere in any other format other than specified above, nor does it apply to quotations, images, artistic works etc that have been reproduced from other sources which may be part of the material to be used.

4. No alteration, omission or addition is made to the material without our written consent. Permission must be re-cleared with Oxford University Press if/when you decide to reprint.

5. The following credit line appears wherever the material is used: author, title, journal, year, volume, issue number, pagination, by permission of Oxford University Press or the sponsoring society if the journal is a society journal. Where a journal is being published on behalf of a learned society, the details of that society must be included in the credit line.

6. For the reproduction of a full article from an Oxford University Press journal for whatever purpose, the corresponding author of the material concerned should be informed of the proposed use. Contact details for the corresponding authors of all Oxford University Press journal contact can be found alongside either the abstract or full text of the article concerned, accessible from www.oxfordjournals.org Should there be a problem clearing these rights, please contact journals.permissions@oup.com

7. If the credit line or acknowledgement in our publication indicates that any of the figures, images or photos was reproduced, drawn or modified from an earlier source it will be necessary for you to clear this permission with the original publisher as well. If this permission has not been obtained, please note that this material cannot be included in your publication/photocopies.

8. While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by Oxford University Press or by Copyright Clearance Center (CCC)) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and Oxford University Press reserves the right to take any and all action to protect its copyright in the materials.

9. This license is personal to you and may not be sublicensed, assigned or transferred by you to any other person without Oxford University Press's written permission.

10. Oxford University Press reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

11. You hereby indemnify and agree to hold harmless Oxford University Press and CCC, and their respective officers, directors, employs and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

12. Other Terms and Conditions:

v1.4

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.